KV6017: Enterprise Networks and Security

GNS3 Assignment

Abdulla Abdulrahman AlBassam

W23055814

BSc (Hons) Networks & Cyber Security

January 15th, 2026

Word Count: 3004

AI Declaration:

I used Claude & ChatGPT to help find relevant sources where necessary and pinpoint specific sections in those sources that are useful for the specific task I had at hand.

I used ChatGPT to organise and cite the references used.

I used ChatGPT to help with grammar and formatting.

I used Claude & ChatGPT to give ideas as to what could be trimmed to help decrease the length of my video submission (From just over 10:00 to the final 8:48 submission)

I used ChatGPT to help organise my "brain dump" into a concise and organised script for the video recording.

I used Claude to ensure my report covered all requirements, allowing me to review the report against my topology configurations.

I used Claude to help me find the technical jargon for non-technical terms I gave it. (E.g. "Propagate").

I used ChatGPT to help consolidate my word count (helped me in cutting out the "fluff').

Section 1: Network Decisions and Justification

1.1 Requirement 1: Secure Direct Link Between EXEC and HR

The EXEC and HR departments require a secure, dedicated communication channel that operates independently from the main MPLS backbone. This requires creating a private tunnel between two sites that already have connectivity through a shared P network.

Technology Selection

The implementation combines GRE tunnelling with IPSec encryption. GRE was selected as the tunnelling protocol because it creates a virtual point-to-point link between EXEC and HR, effectively making them appear as directly connected neighbours regardless of the underlying network topology. Unlike IPSec tunnel mode alone, GRE supports multicast traffic and routing protocols, which provides flexibility for future expansion (Cisco Systems, n.d. -g).

However, GRE alone provides no security, it just encapsulates packets. This is why IPSec is used. By wrapping the GRE tunnel inside IPSec, all traffic between EXEC and HR becomes encrypted. The configuration uses IKEv1 for the initial security association establishment, with AES-256 for encryption and SHA for integrity verification.

Implementation Decisions

AES-256 was chosen over alternatives like 3DES because it offers superior security with better performance on modern hardware. The 256-bit key length exceeds current compliance requirements and provides headroom against future cryptographic advances (Cisco Systems, n.d.-e). SHA-HMAC ensures packet integrity, preventing any tampering during transit (Cisco Systems, n.d.-f).

The tunnel endpoints use dedicated IP addresses (192.168.100.1 and 192.168.100.2) on a /30 subnet, which is standard practice for point-to-point links. This addressing scheme clearly identifies tunnel traffic in routing tables and simplifies troubleshooting. The physical transport runs over the existing WAN links to the MAN-HO PE router, meaning no additional infrastructure was required since the security layer sits entirely in software.

This design ensures that if an attacker compromised the provider network, EXEC-HR communications would remain confidential. The separation from the MPLS VPN also means this traffic flow has its own failure domain, issues with the broader VPN infrastructure would not affect this link.

1.2 Requirement 2: Full-Mesh MPLS VPN for HR, PROD-1, and LOGISTICS

Three geographically dispersed sites: HR in Manchester, PROD-1 in Leeds, and LOGISTICS in Newcastle, require seamless any-to-any connectivity. Each site must communicate directly with the others without traffic backhauling through a central hub. This is a textbook use case for MPLS Layer 3 VPN technology.

Technology Selection

MPLS was selected as the transport technology because it enables traffic engineering and QoS capabilities that traditional IP routing cannot match. More importantly, MPLS supports Virtual Routing and Forwarding (VRF) instances, which create logically separate routing tables on shared infrastructure.

The VRF named "VRF-MESH" isolates these three sites from other network traffic. Each site's routes exist only within this VRF, providing inherent security through separation. A user at LOGISTICS cannot accidentally (or intentionally) route traffic to networks outside their permitted controls (Cisco Systems, n.d.-h).

Implementation Decisions

BGP with VPNv4 address family handles route distribution between Provider Edge (PE) routers. The Route Distinguisher (RD) 100:1 uniquely identifies VRF-MESH routes within the BGP process, while Route Targets control import/export policies. By configuring both import and export of RT 100:1 on all three PE-CE connections, routes automatically propagate to create the full mesh, no manual route configuration is required at each site (Cisco Systems, n.d.-i).

Label Distribution Protocol (LDP) establishes the MPLS label-switched paths across the backbone. The P (Provider) router and all PE routers participate in LDP neighbour relationships, ensuring that labelled paths exist between any two points. This label stack (an inner VPN label + an outer transport label) allows the core network to forward traffic without examining the customer IP addresses.

The whole point of this design is its scalability. Adding a fourth site to the mesh requires only configuring the new PE-CE connection with appropriate RT values; BGP automatically distributes the routes to all existing participants. Contrast this with a traditional hub-and-spoke model, where adding sites requires touching configurations at the hub, MPLS VPN scales far more effortlessly.

1.3 Requirement 3: MPLS VPN Between EXEC and R-and-D with Selective Encryption

The R-and-D department involves two distinct network segments: LAN A for sensitive data, and LAN B containing non-sensitive data. Both require connectivity to EXEC, but only LAN A traffic requires encryption.

Technology Selection

A second VRF, named "VRF-EXEC-RD" with RD 100:2, provides the routing infrastructure. Keeping EXEC and R-and-D in a separate VRF from the HR-PROD-1-LOGISTICS mesh ensures complete traffic isolation between these two groups. The MPLS backbone carries both VRFs simultaneously, but their routing tables never intersect.

For the encryption requirement, IPSec with crypto access lists provides specific control. Rather than encrypting all traffic (which would waste CPU power on non-sensitive data), the access list explicitly matches only LAN A's subnet (10.3.1.0/26). Traffic from LAN B (10.3.2.0/28) passes through the MPLS VPN unencrypted (Cisco Systems, n.d.-j).

Implementation Decisions

The crypto map applied to the tunnel interface references an access list that permits only 10.3.1.0/26 source addresses. This approach means the router evaluates each packet against the ACL, matching packets entering the IPSec tunnel while non-matching packets route normally via MPLS. Unlike requirement 1, this implementation uses AES-128 rather than AES-256. The brief specifies that only low sensitivity data traverses this link and that security measures should be weighed against performance penalties. AES-128 provides strong encryption while reducing computational overhead, resulting in lower latency and higher throughput which is an appropriate trade-off given the data’s classification (Cisco Systems, n.d.-e).

This design demonstrates that security does not need to be all-or-nothing. By layering IPSec selectively atop the MPLS VPN, the network delivers appropriate protection levels based on data sensitivity, a principle that aligns with modern zero-trust architectures where access controls are as granular as a user’s requirements demand.

1.4 Requirement 4: SQL Server Placement and Access Control

The SQL server must be placed in either PROD-1 (LDS) or LOGISTICS (NCL), with the decision based on performance modelling for HR users at MAN_HO. Additionally, access must be restricted so that only HR can reach the server, and only when sessions originate from the HR LAN.

Performance Analysis

Transfer time consists of propagation delay + transmission time. Propagation delay is calculated as distance divided by the speed of light in fibre (approximately 200,000 km/s). Transmission time equals file size divided by bandwidth. For LDS (PROD-1): Propagation delay = 80 km ÷ 200,000 km/s = 0.4 ms. Transmission time for 100MB = (100 × 8 Mb) ÷ 70 Mbps = 11.43 seconds. Total transfer time (TTT) = 11.43 seconds. For NCL (LOGISTICS): Propagation delay = 220 km ÷ 200,000 km/s = 1.1 ms. Transmission time for 100MB = (100 × 8 Mb) ÷ 100 Mbps = 8.0 seconds. TTT = 8.0 seconds. Despite the greater distance, LOGISTICS provides 30% faster transfer times due to higher bandwidth. The propagation delay difference (0.7 ms) is negligible compared to the 3.43s transmission time improvement. Therefore, the SQL server is placed at LOGISTICS.

Technology Selection

Extended Access Control Lists (ACLs) provide packet filtering based on source address, destination address, and port number. While firewalls offer more sophisticated inspection capabilities, ACLs are native to Cisco IOS, require no additional hardware, and execute at line rate. For this requirement: permit HR, deny everyone else, with session-origination enforcement, an extended ACL with the established keyword is the appropriate tool (Cisco Systems, n.d.-k).

Implementation Decisions

The ACL "SQL-SERVER-ACCESS" uses a specific permit-deny structure. The first entry allows TCP traffic from HR's subnet (10.1.2.0/25) to the SQL server on port 1433, the standard Microsoft SQL Server port. A subsequent entry permits return traffic using the established keyword, which checks for ACK or RST flags which ensures responses are only allowed for sessions initiated by HR, satisfying the session-origination requirement (Cisco

Systems, n.d.-l). The third entry explicitly denies all other traffic to the SQL server with logging enabled to track access attempts. The final entry permits all other traffic, ensuring the ACL does not inadvertently block legitimate LOGISTICS LAN communications. This approach (specific permits, specific denies, general permit) follows Cisco best practices for ACL design. The ACL is applied inbound on the LOGISTICS CE router's interface facing the PE, filtering traffic before it enters the local network.

1.5 Requirement 5: PROD-2 Isolation from EXEC

PROD-2 at the Newcastle site must not communicate with EXEC under any circumstances. This isolation requirement is satisfied by the existing VRF architecture.

PROD-2 connects to NCL's global routing table (no VRF assignment), while EXEC belongs to VRF-EXEC-RD (RD 100:2). Since VRFs maintain completely separate routing tables from the global table, no route to EXEC exists within PROD-2's routing context and vice versa (Cisco Systems, n.d.-m). Packets from PROD-2 destined for EXEC's subnet have no forwarding path; the router returns "destination unreachable" without any explicit deny configs. This separation-by-design is more robust than ACL blocking because it operates at the routing plane rather than requiring packet-by-packet inspection.

1.6 Requirement 6: DHCP Services Across All LANs

Dynamic Host Configuration Protocol (DHCP) automates IP address assignment across all customer LANs, reducing administrative overhead and eliminating manual configuration errors on end-user devices.

Technology Selection

Each CE router runs a local DHCP server for its connected LAN. Centralising DHCP on a single server would require DHCP relay configuration and introduce a single point of failure. The distributed model ensures that each site operates independently, and a WAN outage would not prevent local devices from obtaining addresses (Cisco Systems, n.d.-n).

Addressing Scheme

The IP addressing plan uses Variable Length Subnet Masking (VLSM) to allocate appropriately sized subnets based on each department's user count. This approach conserves address space while allowing room for growth. Table 1 presents the complete addressing scheme.

Table 1

IP Addressing Scheme and DHCP Configuration

Site	Subnet	Mask	Users	Gateway	DHCP Pool
EXEC	10.1.1.0	/26	50	10.1.1.1	EXEC-LAN
HR	10.1.2.0	/25	80	10.1.2.1	HR-LAN
PROD-1	10.2.1.0	/25	120	10.2.1.1	PROD1-LAN
R-and-D LAN A	10.3.1.0	/26	50	10.3.1.1	RD-LAN-A
R-and-D LAN B	10.3.2.0	/28	12	10.3.2.1	RD-LAN-B
PROD-2	10.3.2.0	/23	300	10.3.3.1	PROD2-LAN
LOGISTICS	10.3.4.0	/25	80	10.3.4.1	LOGISTICS-LAN

Each DHCP pool excludes the gateway address and reserves the first few addresses for infrastructure devices. The default gateway and DNS server parameters are automatically distributed to clients, ensuring zero-touch deployment for end-user workstations. This configuration supports the specified user counts while maintaining approximately 20% room for growth.

Section 2: Router Configurations and Testing

2.1 Requirement 1: Secure Direct Link Between EXEC and HR (GRE over IPSec)

Router Configurations

EXEC Router - IPSec and GRE Tunnel Config:

```
crypto isakmp policy 10  - Security rules for setting up the encrypted connection
encr aes 256  - 256 bit encryption
authentication pre-share – routers to share passwords
group 5 - Diffie-Hellman, securely exchanges encryptions keys between routers

crypto isakmp key 1234 address 10.1.100.2 - Shared password to connect with HR router

crypto ipsec transform-set EXEC-HR-TRANSFORM esp-aes 256 esp-sha-hmac – defines how real traffic will be encrypted

crypto map EXEC-HR-CRYPTOMAP 10 ipsec-isakmp - Crypto map definition,links encryption settings together
set peer 10.1.100.2 - IPSec peer is HR router
set transform-set EXEC-HR-TRANSFORM – applies encryption method defined above
match address EXEC-HR-CRYPTO-ACL – tells router which traffic to encrypt

ip access-list extended EXEC-HR-CRYPTO-ACL – Access list specifying what gets encrypted
permit gre host 10.1.100.1 host 10.1.100.2  - Encrypts GRE tunnel traffic

interface Tunnel0 – Creates virtual tunnel interface
ip address 192.168.100.1 255.255.255.252  - Tunnel endpoint address
tunnel source FastEthernet1/0  - Physical source interface
tunnel destination 10.1.100.2  - Tunnel destination (HR)

interface FastEthernet1/0  - Physical interface to HR
ip address 10.1.100.1 255.255.255.252  - Direct link address
crypto map EXEC-HR-CRYPTOMAP  - Apply IPSec encryption

ip route 10.1.2.0 255.255.255.128 Tunnel0  - Route HR LAN via tunnel
```

HR Router - IPSec and GRE Tunnel Config:

```
crypto isakmp policy 10  - Matching IKE policy
encr aes 256  - Matches EXEC
authentication pre-share – Same as above
group 5 – Same as above

crypto isakmp key 1234 address 10.1.100.1 - Key for EXEC peer

crypto ipsec transform-set HR-EXEC-TRANSFORM esp-aes 256 esp-sha-hmac  - Matching transform

crypto map HR-EXEC-CRYPTOMAP 10 ipsec-isakmp – Same as above
```

set peer 10.1.100.1 *- Peer is EXEC*
set transform-set HR-EXEC-TRANSFORM – *Same as above*
match address HR-EXEC-CRYPTO-ACL – *Same as above*

ip access-list extended HR-EXEC-CRYPTO-ACL – *Same as above*
permit gre host 10.1.100.2 host 10.1.100.1 *- Encrypts GRE to EXEC*

interface Tunnel0 – *Same as above*
ip address 192.168.100.2 255.255.255.252 *- Other tunnel endpoint*
tunnel source FastEthernet1/0 – *Same as above*
tunnel destination 10.1.100.1 – *Same as above*

interface FastEthernet1/0 – *Same as above*
ip address 10.1.100.2 255.255.255.252 – *Same as above*
crypto map HR-EXEC-CRYPTOMAP *- Applies crypto map*

ip route 10.1.1.0 255.255.255.192 Tunnel0 *- Route EXEC LAN via tunnel*

Testing Evidence

Test 1.1: Running Configuration - IPSec on EXEC

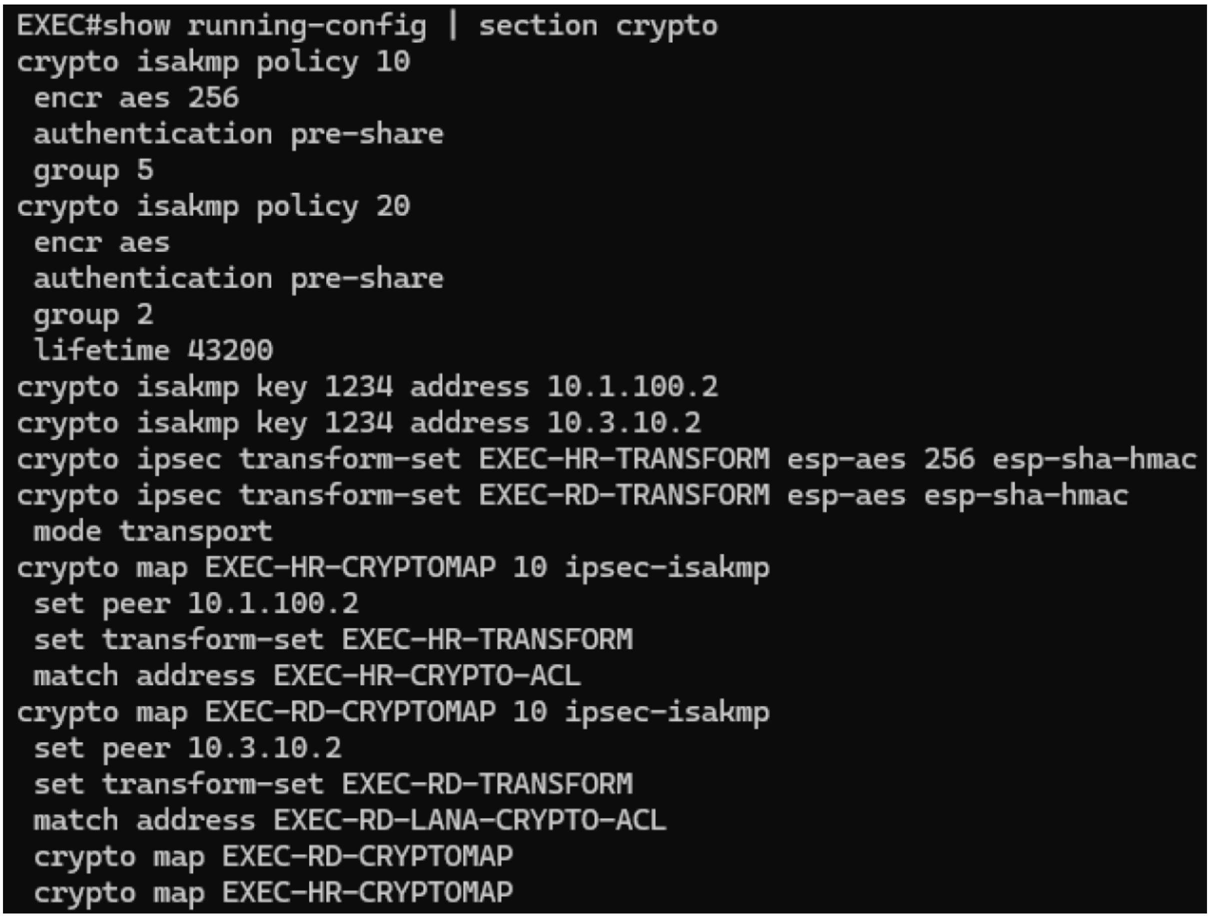

```
EXEC#show running-config | section crypto
crypto isakmp policy 10
 encr aes 256
 authentication pre-share
 group 5
crypto isakmp policy 20
 encr aes
 authentication pre-share
 group 2
 lifetime 43200
crypto isakmp key 1234 address 10.1.100.2
crypto isakmp key 1234 address 10.3.10.2
crypto ipsec transform-set EXEC-HR-TRANSFORM esp-aes 256 esp-sha-hmac
crypto ipsec transform-set EXEC-RD-TRANSFORM esp-aes esp-sha-hmac
 mode transport
crypto map EXEC-HR-CRYPTOMAP 10 ipsec-isakmp
 set peer 10.1.100.2
 set transform-set EXEC-HR-TRANSFORM
 match address EXEC-HR-CRYPTO-ACL
crypto map EXEC-RD-CRYPTOMAP 10 ipsec-isakmp
 set peer 10.3.10.2
 set transform-set EXEC-RD-TRANSFORM
 match address EXEC-RD-LANA-CRYPTO-ACL
 crypto map EXEC-RD-CRYPTOMAP
 crypto map EXEC-HR-CRYPTOMAP
```

Explanation: This shows the IPSec configuration including ISAKMP policy, transform set, and crypto map applied to the interface.

Test 1.2: Running Configuration - Tunnel on EXEC

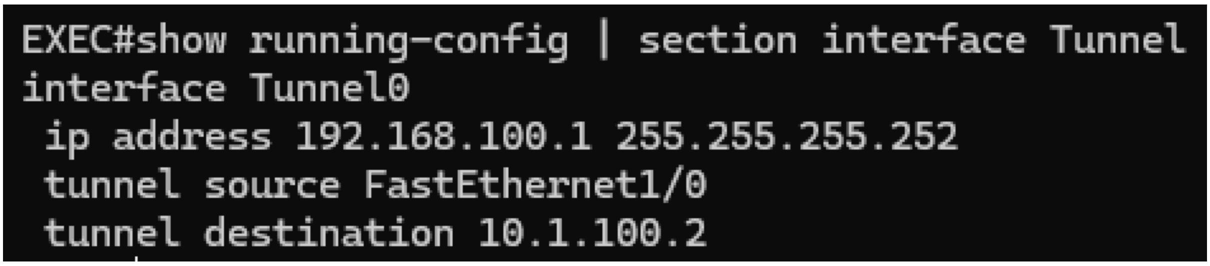

```
EXEC#show running-config | section interface Tunnel
interface Tunnel0
 ip address 192.168.100.1 255.255.255.252
 tunnel source FastEthernet1/0
 tunnel destination 10.1.100.2
```

Explanation: This shows the GRE tunnel configuration with source, destination, and IP address.

Test 1.3: IKE Security Association

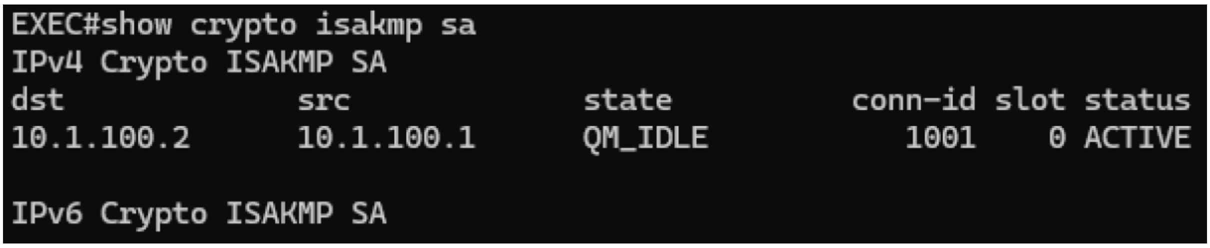

```
EXEC#show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst             src             state          conn-id slot status
10.1.100.2      10.1.100.1      QM_IDLE           1001    0 ACTIVE

IPv6 Crypto ISAKMP SA
```

Explanation: The QM_IDLE state confirms IKE Phase 1 completed successfully. The connection shows between 10.1.100.1 (EXEC) and 10.1.100.2 (HR).

Test 1.4: IPSec Security Association

Router: EXEC / **Command:** show crypto ipsec sa

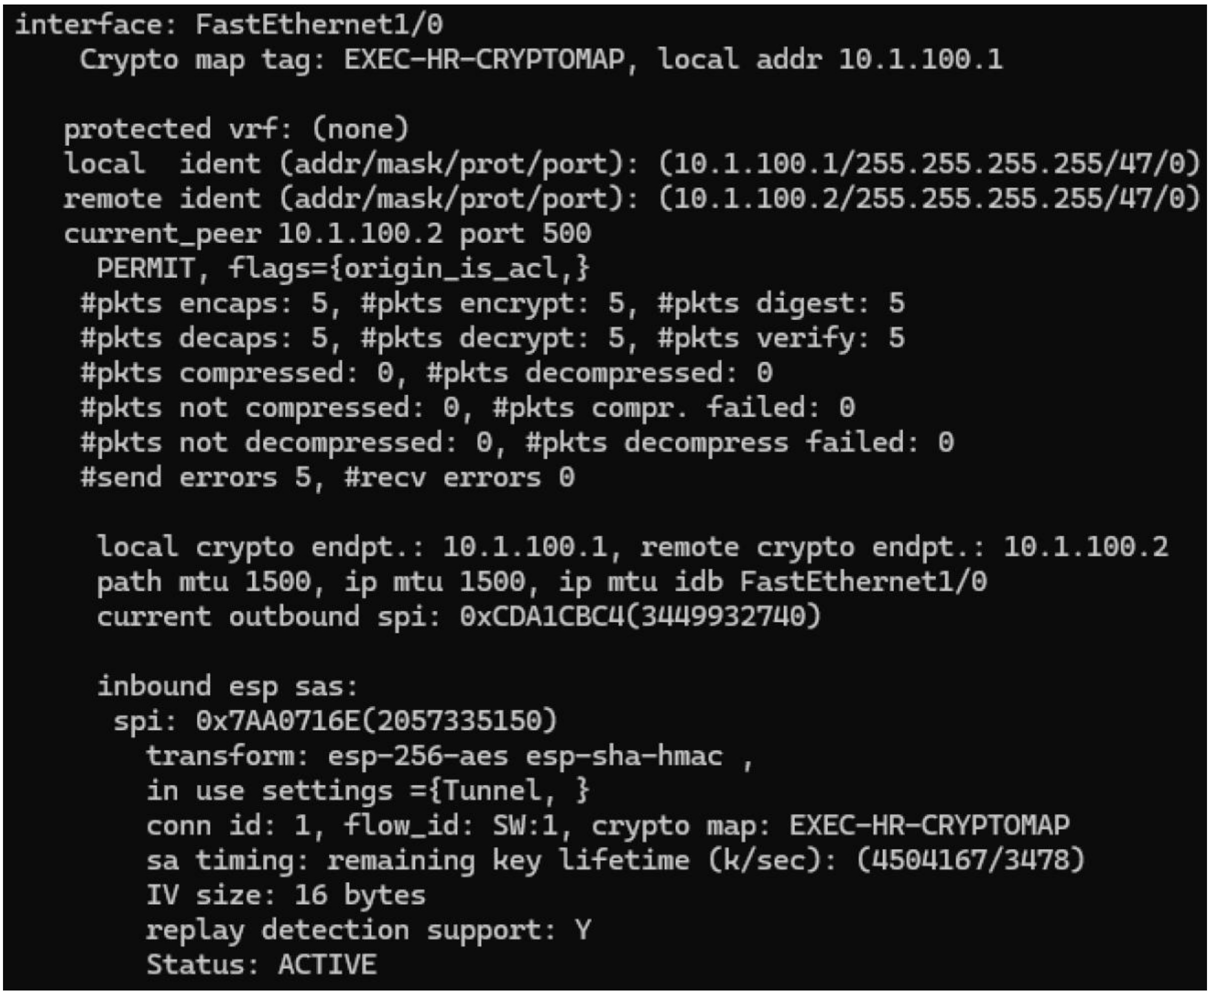

```
interface: FastEthernet1/0
    Crypto map tag: EXEC-HR-CRYPTOMAP, local addr 10.1.100.1

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (10.1.100.1/255.255.255.255/47/0)
   remote ident (addr/mask/prot/port): (10.1.100.2/255.255.255.255/47/0)
   current_peer 10.1.100.2 port 500
     PERMIT, flags={origin_is_acl,}
    #pkts encaps: 5, #pkts encrypt: 5, #pkts digest: 5
    #pkts decaps: 5, #pkts decrypt: 5, #pkts verify: 5
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 5, #recv errors 0

     local crypto endpt.: 10.1.100.1, remote crypto endpt.: 10.1.100.2
     path mtu 1500, ip mtu 1500, ip mtu idb FastEthernet1/0
     current outbound spi: 0xCDA1CBC4(3449932740)

     inbound esp sas:
      spi: 0x7AA0716E(2057335150)
        transform: esp-256-aes esp-sha-hmac ,
        in use settings ={Tunnel, }
        conn id: 1, flow_id: SW:1, crypto map: EXEC-HR-CRYPTOMAP
        sa timing: remaining key lifetime (k/sec): (4504167/3478)
        IV size: 16 bytes
        replay detection support: Y
        Status: ACTIVE
```

Explanation: Shows AES-256 encryption, SHA integrity, and packet counters. Non-zero encrypt/decrypt counts prove encrypted traffic is flowing.

Test 1.5: Tunnel Interface Status

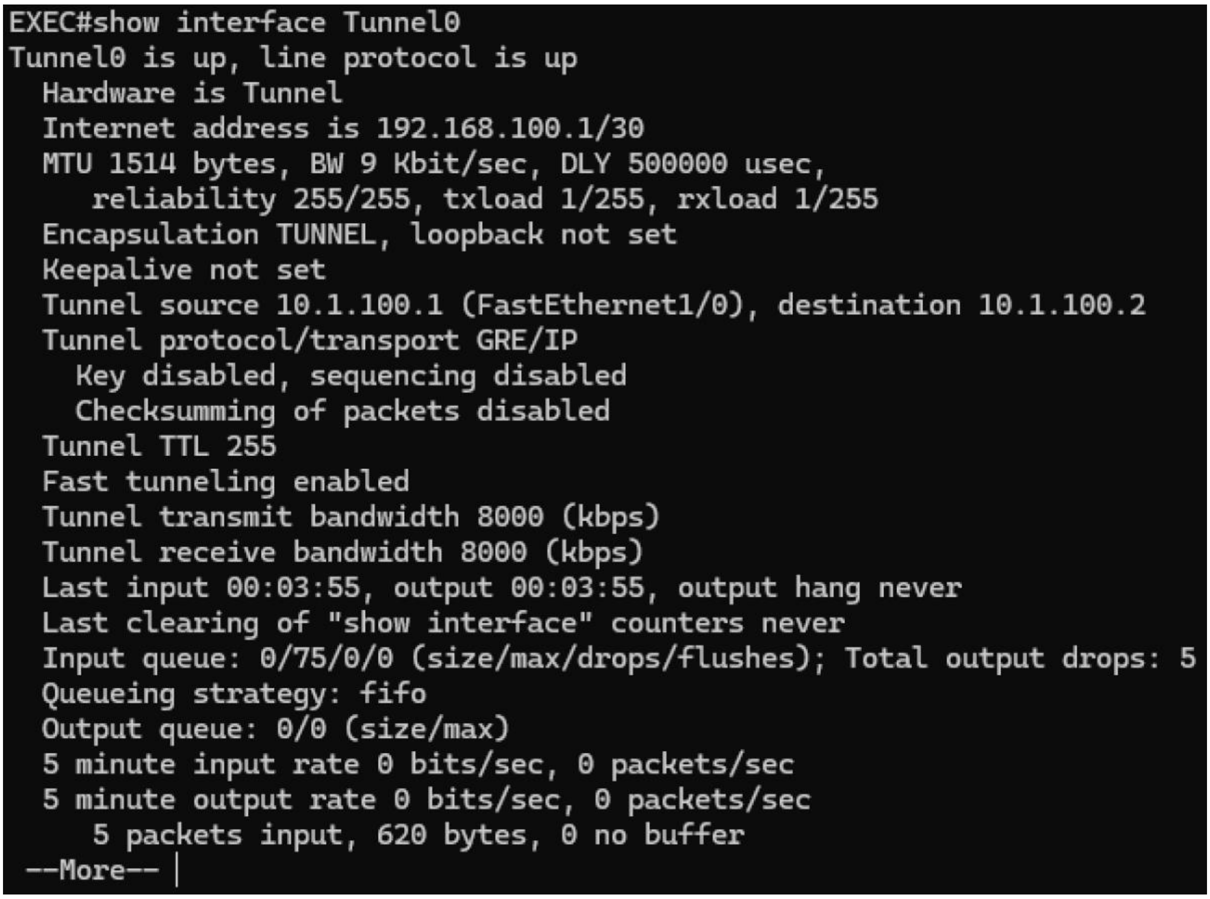

```
EXEC#show interface Tunnel0
Tunnel0 is up, line protocol is up
  Hardware is Tunnel
  Internet address is 192.168.100.1/30
  MTU 1514 bytes, BW 9 Kbit/sec, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 10.1.100.1 (FastEthernet1/0), destination 10.1.100.2
  Tunnel protocol/transport GRE/IP
    Key disabled, sequencing disabled
    Checksumming of packets disabled
  Tunnel TTL 255
  Fast tunneling enabled
  Tunnel transmit bandwidth 8000 (kbps)
  Tunnel receive bandwidth 8000 (kbps)
  Last input 00:03:55, output 00:03:55, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 5
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     5 packets input, 620 bytes, 0 no buffer
 --More--
```

Explanation: Tunnel0 shows up/up status confirming GRE tunnel is operational.

Test 1.6: Ping Through Tunnel

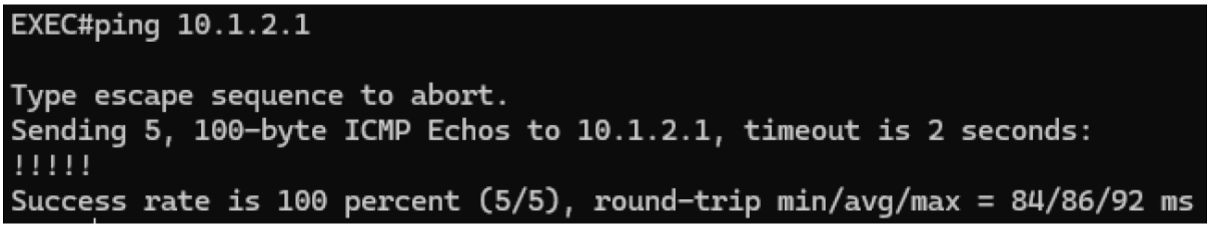

```
EXEC#ping 10.1.2.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.2.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 84/86/92 ms
```

Explanation: 100% success rate confirms connectivity from EXEC to HR LAN through the encrypted tunnel.

Test 1.7: Traceroute Through Tunnel

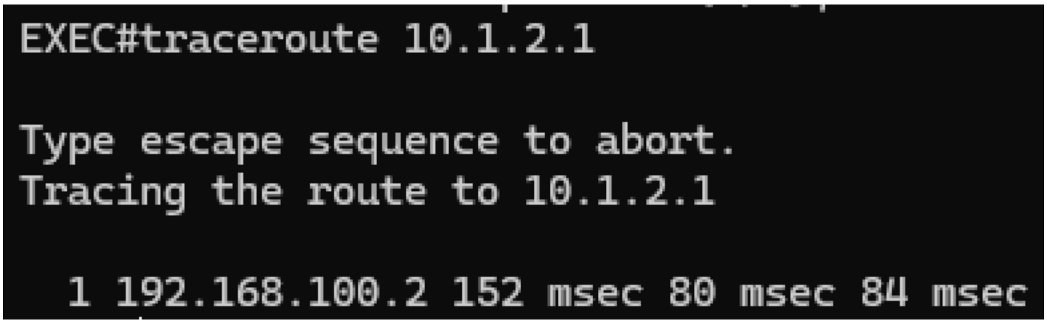

```
EXEC#traceroute 10.1.2.1

Type escape sequence to abort.
Tracing the route to 10.1.2.1

  1 192.168.100.2 152 msec 80 msec 84 msec
```

Explanation: Shows traffic path via tunnel (192.168.100.2 addresses), proving the dedicated secure link is used rather than MPLS backbone.

2.2 Requirement 2: Full-Mesh MPLS VPN (HR, PROD-1, LOGISTICS)

Router Configurations

P Router (Provider Core) - MPLS Backbone:

```
interface Loopback0 – Unique Router ID for MPLS/BGP
ip address 1.1.1.1 255.255.255.255

interface FastEthernet0/0  - To MAN-HO PE
ip address 172.20.20.1 255.255.255.252
mpls ip - Enables MPLS label switching

interface FastEthernet0/1  - To LDS PE
ip address 172.20.20.5 255.255.255.252
mpls ip - Enables MPLS label switching

interface FastEthernet1/0  - To NCL PE
ip address 172.20.20.9 255.255.255.252
mpls ip - Enables MPLS label switching

router ospf 1  - IGP for loopback reachability, enables OSPF so routers can find each other
network 1.1.1.1 0.0.0.0 area 0 – Advertises P routers loopback
network 172.20.20.0 0.0.0.3 area 0 – Advertises link to MAN_HO
network 172.20.20.4 0.0.0.3 area 0 – Link to LDS
network 172.20.20.8 0.0.0.3 area 0 – Link to NCL

mpls ldp router-id Loopback0 force  - Uses loopback for LDP
```

MAN-HO Router (PE) - VRF-MESH Configuration:

```
ip vrf VRF-MESH  - Creates isolated routing table for HR, Prod-1 and LOGISTICS
rd 100:1  - Route ID for this VRF
route-target export 100:1  - Tags outgoing routers so other sites can receive them
route-target import 100:1  - Accepts routes tagged with 100:1

interface Loopback0 – Virtual interface incase physical link fails
ip address 2.2.2.2 255.255.255.255 – router ID for MAN-HO

interface FastEthernet1/0  - To HR CE
ip vrf forwarding VRF-MESH  - Assigns this route to VRF
ip address 10.1.10.5 255.255.255.252 – Link address to HR

interface FastEthernet2/0  - To P router
ip address 172.20.20.2 255.255.255.252 – Link address to MPLS backbone
mpls ip – Same as above

router bgp 65000 – Starts BGP process for route exchange
neighbor 3.3.3.3 remote-as 65000  - BGP connection to LDS router for route sharing
neighbor 3.3.3.3 update-source Loopback0 – Loopback for stable BGP
neighbor 4.4.4.4 remote-as 65000  - BGP to NCL
neighbor 4.4.4.4 update-source Loopback0 – Loopback for stable BGP
```

```
 address-family vpnv4 – Enables VPN route exchange between PE routers
 neighbor 3.3.3.3 activate – Turns on VPNv4 with LDS
 neighbor 3.3.3.3 send-community both  - Includes route target tags so VPN routes are shared
  neighbor 4.4.4.4 activate – Turns on VPNv4 with NCL
  neighbor 4.4.4.4 send-community both – Same as above

 address-family ipv4 vrf VRF-MESH – BGP settings specific to VRF mesh
 redistribute connected – shares directly connected networks into BGP
 redistribute static – shares static routes into BGP

ip route vrf VRF-MESH 10.1.2.0 255.255.255.128 10.1.10.6  - To HR LAN

mpls ldp router-id Loopback0 force – Uses loopback address fo rthe MPLS label distribution
```

LDS Router (PE) - VRF-MESH for PROD-1:

```
ip vrf VRF-MESH  - Same VRF definition
rd 100:1  - Same RD
route-target export 100:1 – Tags outgoing routes
route-target import 100:1 – Accepts incoming routes

interface Loopback0 – Same as above
ip address 3.3.3.3 255.255.255.255 – Router ID for LDS

interface FastEthernet0/1  - To PROD-1 CE
ip vrf forwarding VRF-MESH – Places this interface in VRF mesh
ip address 10.2.10.1 255.255.255.252 – Link address to PROD-1

ip route vrf VRF-MESH 10.2.1.0 255.255.255.128 10.2.10.2  - To PROD-1 LAN
```

NCL Router (PE) - VRF-MESH for LOGISTICS:

```
ip vrf VRF-MESH – Same VRF def as other PE routers
rd 100:1 – Same distinguisher
route-target export 100:1 – Same as above
route-target import 100:1 – Same as above

interface Loopback0 – Same as above
ip address 4.4.4.4 255.255.255.255 Router ID for NCL

interface FastEthernet2/0  - To LOGISTICS CE
ip vrf forwarding VRF-MESH – Same as above
ip address 10.3.10.9 255.255.255.252 – Link address to LOGISTICS

ip route vrf VRF-MESH 10.3.4.0 255.255.255.128 10.3.10.10  - To LOGISTICS LAN
```

Testing Evidence

Test 2.1: MPLS Forwarding Table

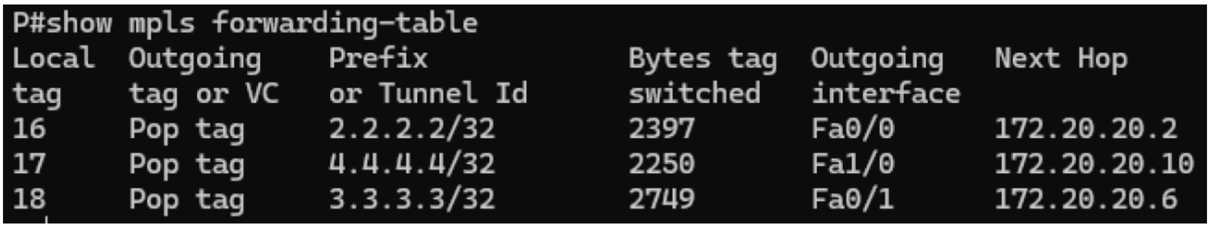

```
P#show mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
16     Pop tag     2.2.2.2/32        2397       Fa0/0      172.20.20.2
17     Pop tag     4.4.4.4/32        2250       Fa1/0      172.20.20.10
18     Pop tag     3.3.3.3/32        2749       Fa0/1      172.20.20.6
```

Explanation: Shows MPLS labels assigned to each PE router loopback. Labels 16, 17, 18 indicate successful LDP label distribution.

Test 2.2: LDP Neighbour Relationships

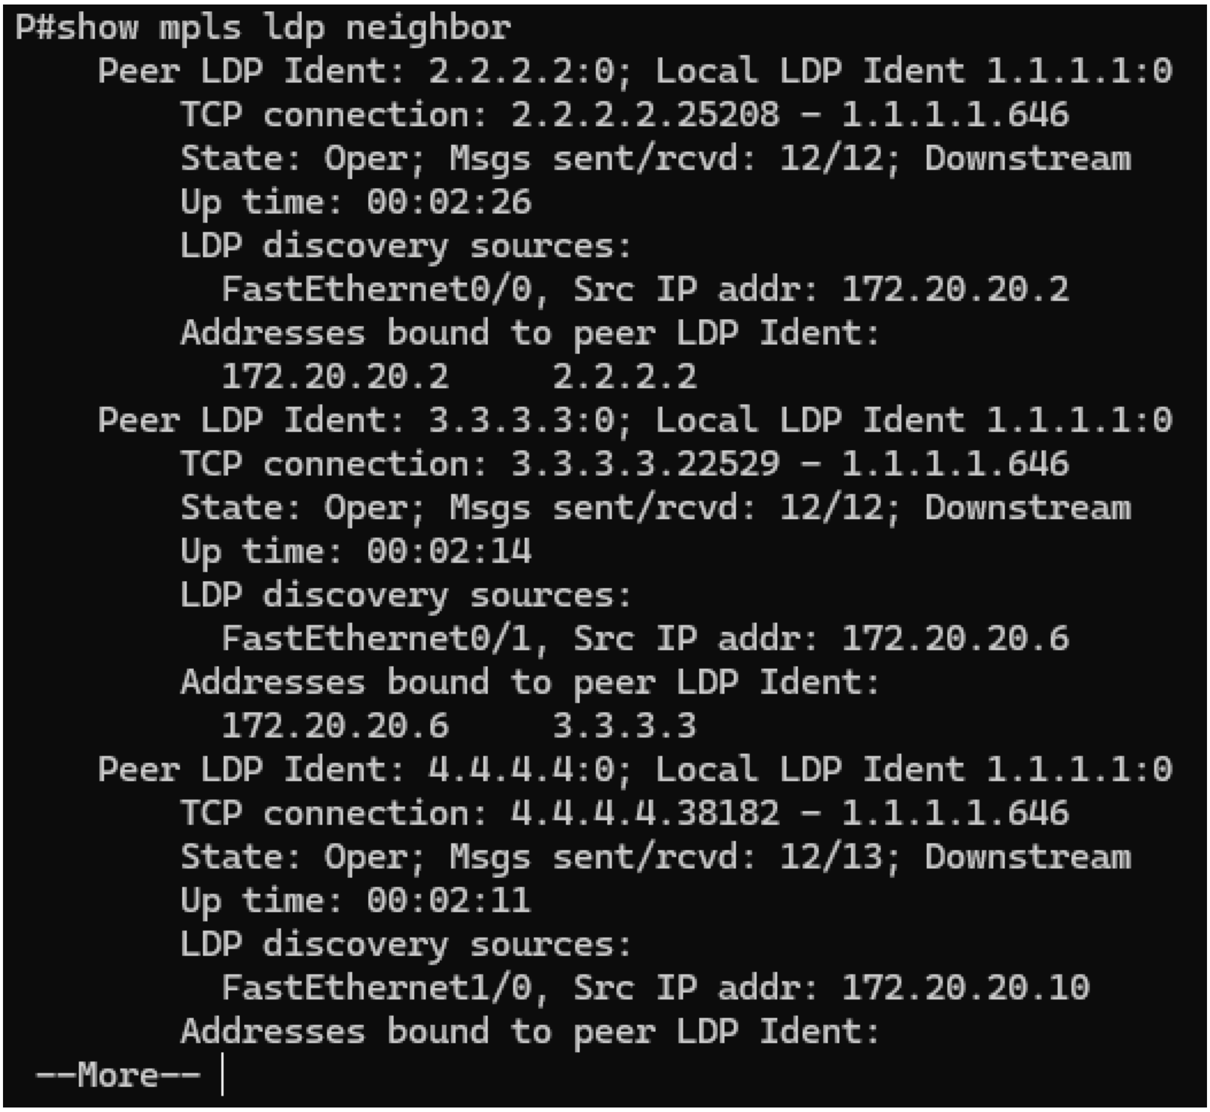

```
P#show mpls ldp neighbor
    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 1.1.1.1:0
        TCP connection: 2.2.2.2.25208 - 1.1.1.1.646
        State: Oper; Msgs sent/rcvd: 12/12; Downstream
        Up time: 00:02:26
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 172.20.20.2
        Addresses bound to peer LDP Ident:
          172.20.20.2     2.2.2.2
    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 1.1.1.1:0
        TCP connection: 3.3.3.3.22529 - 1.1.1.1.646
        State: Oper; Msgs sent/rcvd: 12/12; Downstream
        Up time: 00:02:14
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 172.20.20.6
        Addresses bound to peer LDP Ident:
          172.20.20.6     3.3.3.3
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 1.1.1.1:0
        TCP connection: 4.4.4.4.38182 - 1.1.1.1.646
        State: Oper; Msgs sent/rcvd: 12/13; Downstream
        Up time: 00:02:11
        LDP discovery sources:
          FastEthernet1/0, Src IP addr: 172.20.20.10
        Addresses bound to peer LDP Ident:
 --More--
```

Explanation: Shows LDP sessions with all three PE routers (2.2.2.2, 3.3.3.3, 4.4.4.4). 'Oper' status confirms operational.

Test 2.3: VRF Routing Table

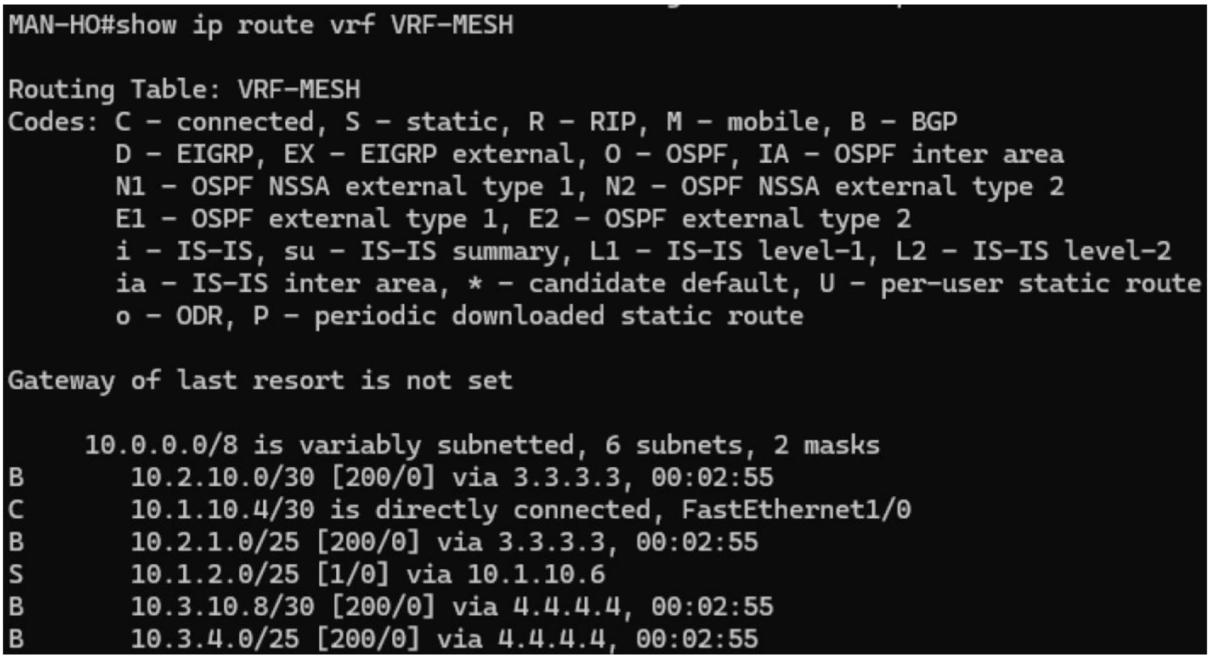

```
MAN-HO#show ip route vrf VRF-MESH

Routing Table: VRF-MESH
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
B       10.2.10.0/30 [200/0] via 3.3.3.3, 00:02:55
C       10.1.10.4/30 is directly connected, FastEthernet1/0
B       10.2.1.0/25 [200/0] via 3.3.3.3, 00:02:55
S       10.1.2.0/25 [1/0] via 10.1.10.6
B       10.3.10.8/30 [200/0] via 4.4.4.4, 00:02:55
B       10.3.4.0/25 [200/0] via 4.4.4.4, 00:02:55
```

Explanation: Shows routes to HR (10.1.2.0/25), PROD-1 (10.2.1.0/25), and LOGISTICS (10.3.4.0/25). BGP-learned routes confirm VPNv4 exchange.

Test 2.4: BGP VPNv4 Routes

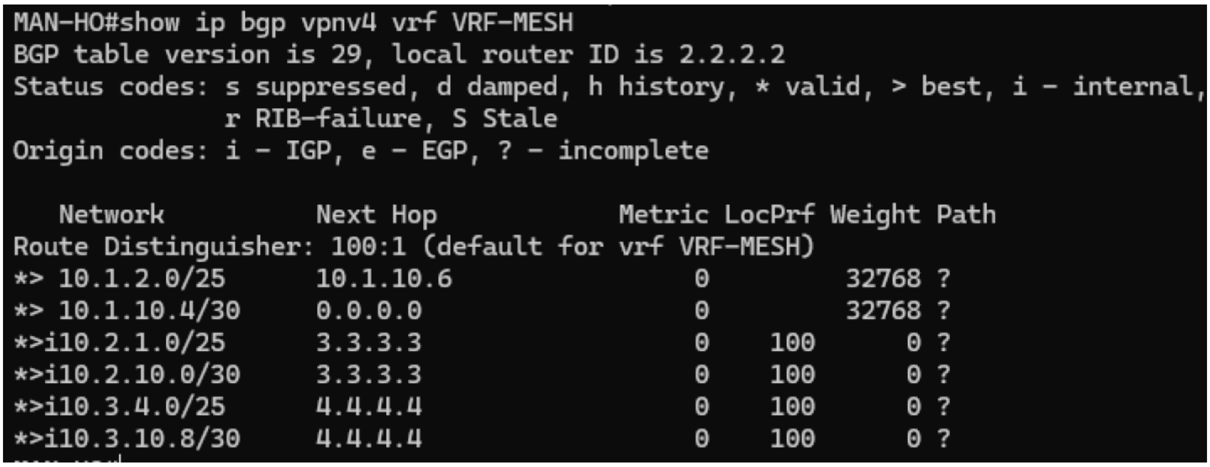

```
MAN-HO#show ip bgp vpnv4 vrf VRF-MESH
BGP table version is 29, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf VRF-MESH)
*> 10.1.2.0/25      10.1.10.6                0         32768 ?
*> 10.1.10.4/30     0.0.0.0                  0         32768 ?
*>i10.2.1.0/25      3.3.3.3                  0    100      0 ?
*>i10.2.10.0/30     3.3.3.3                  0    100      0 ?
*>i10.3.4.0/25      4.4.4.4                  0    100      0 ?
*>i10.3.10.8/30     4.4.4.4                  0    100      0 ?
```

Explanation: Shows all VRF-MESH routes with RD 100:1 and next-hop PE addresses.

Test 2.5: Mesh Connectivity - HR to PROD-1

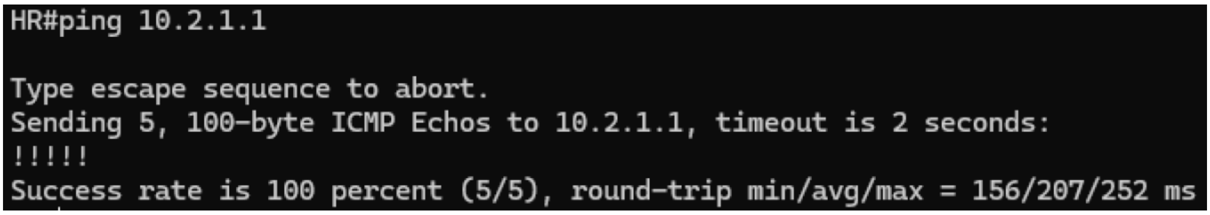

```
HR#ping 10.2.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 156/207/252 ms
```

Explanation: 100% success confirms HR can reach PROD-1 through MPLS VPN.

Test 2.6: Mesh Connectivity - HR to LOGISTICS

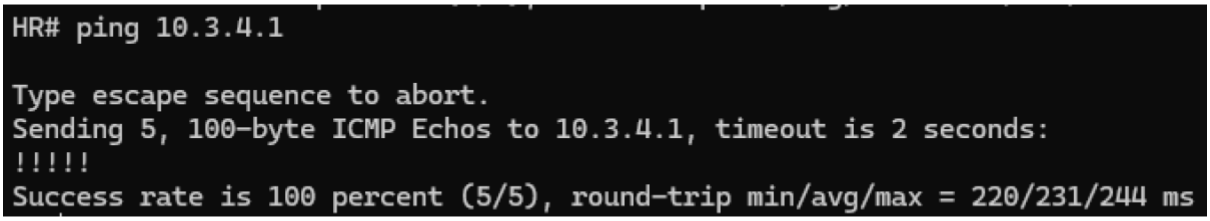

```
HR# ping 10.3.4.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.4.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 220/231/244 ms
```

Explanation: 100% success confirms HR can reach LOGISTICS.

Test 2.7: Mesh Connectivity - PROD-1 to HR

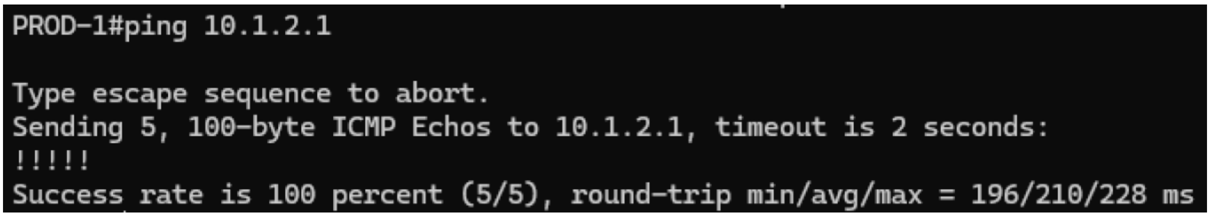

```
PROD-1#ping 10.1.2.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.2.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 196/210/228 ms
```

Explanation: 100% success confirms PROD-1 can reach HR.

Test 2.8: Mesh Connectivity - PROD-1 to LOGISTICS

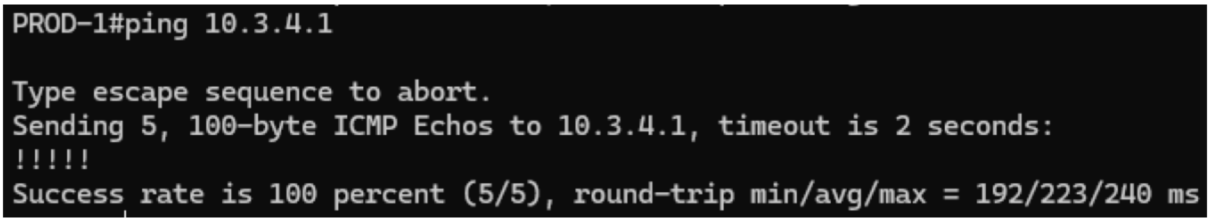

```
PROD-1#ping 10.3.4.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.4.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 192/223/240 ms
```

Explanation: 100% success confirms PROD-1 can reach LOGISTICS.

Test 2.9: Mesh Connectivity - LOGISTICS to HR

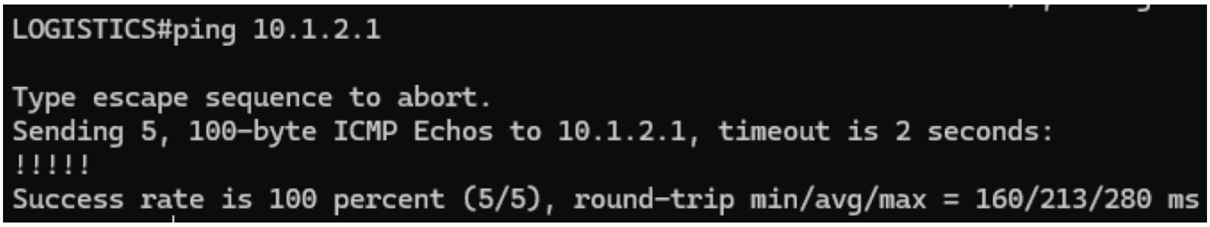

```
LOGISTICS#ping 10.1.2.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.2.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 160/213/280 ms
```

Explanation: 100% success confirms LOGISTICS can reach HR.

Test 2.10: Mesh Connectivity - LOGISTICS to PROD-1

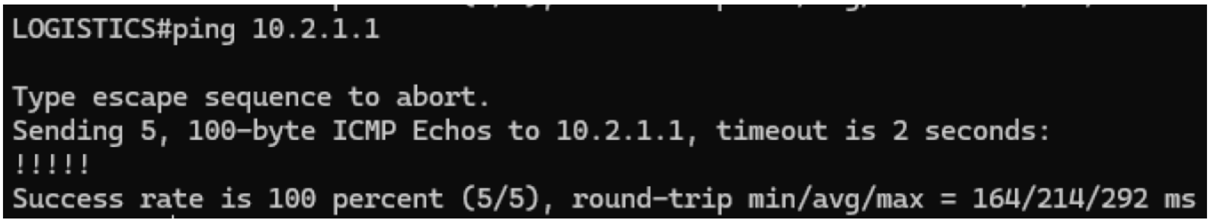

```
LOGISTICS#ping 10.2.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 164/214/292 ms
```

Explanation: 100% success confirms LOGISTICS can reach PROD-1. Full mesh verified.

2.3 Requirement 3: EXEC to R-and-D MPLS VPN with Selective IPSec

Router Configurations

MAN-HO Router (PE) - VRF-EXEC-RD:

```
ip vrf VRF-EXEC-RD  - Creates isolated routing table for EXEC and R and D
rd 100:2  - Different identifier from VRF-MESH
route-target export 100:2 – Same as above
route-target import 100:2 – Same as above

interface FastEthernet0/1  - To EXEC CE
ip vrf forwarding VRF-EXEC-RD – Same as above
ip address 10.1.10.1 255.255.255.252 – Same as above

address-family ipv4 vrf VRF-EXEC-RD  - Enables BGP route sharing for this VRF
redistribute connected – Same as above
redistribute static – Same as above

ip route vrf VRF-EXEC-RD 10.1.1.0 255.255.255.192 10.1.10.2  - To EXEC LAN
```

NCL Router (PE) - VRF-EXEC-RD:

```
ip vrf VRF-EXEC-RD  - Same VRF on remote PE, matches VRF def on NCL router
rd 100:2 – Same as above
route-target export 100:2 – Same as above
route-target import 100:2 – Same as above

interface FastEthernet0/1  - To R-AND-D CE
ip vrf forwarding VRF-EXEC-RD – Same as above
ip address 10.3.10.1 255.255.255.252 – Same as above

ip route vrf VRF-EXEC-RD 10.3.1.0 255.255.255.192 10.3.10.2  - To R-and-D LAN A
ip route vrf VRF-EXEC-RD 10.3.2.0 255.255.255.240 10.3.10.2  - To R-and-D LAN B
```

EXEC Router - Selective IPSec (LAN A Only):

```
crypto isakmp policy 20  - IKE policy for R and D, security rules for encrypted connection
encr aes  - AES-128, lighter than 256
authentication pre-share
group 2
lifetime 43200  - 12-hour lifetime, refreshes every 12 hrs

crypto isakmp key 1234 address 10.3.10.2  - Shared key for R and D

crypto ipsec transform-set EXEC-RD-TRANSFORM esp-aes esp-sha-hmac – Defines AES-128 with SHA integrity checking
mode transport – Encrypts only data not IP headers

crypto map EXEC-RD-CRYPTOMAP 10 ipsec-isakmp – creates encryption policy for R and D traffic
set peer 10.3.10.2 – Specifies R and D router as the encryption partner
set transform-set EXEC-RD-TRANSFORM - Applies AES 128 to settings defined above
```

```
match address EXEC-RD-LANA-CRYPTO-ACL  - Only LAN A traffic encrypted

ip access-list extended EXEC-RD-LANA-CRYPTO-ACL
permit ip 10.1.1.0 0.0.0.63 10.3.1.0 0.0.0.63  - ONLY encrypt EXEC to LAN A traffic

interface FastEthernet0/1  - To MAN-HO PE
crypto map EXEC-RD-CRYPTOMAP  - Apply crypto (encryption for this interface)
```

R-AND-D Router - Matching IPSec:

```
crypto isakmp policy 20 – Security rules matching exec’s policy
encr aes - Same as above
authentication pre-share - Same as above
group 2 - exchange method matching exec
lifetime 43200 – refreshes every 12 hrs

crypto isakmp key 1234 address 10.1.10.2  - Key for EXEC

crypto ipsec transform-set RD-EXEC-TRANSFORM esp-aes esp-sha-hmac
mode transport

ip access-list extended RD-EXEC-LANA-CRYPTO-ACL - Same as above
permit ip 10.3.1.0 0.0.0.63 10.1.1.0 0.0.0.63  - Mirror of EXEC ACL

interface FastEthernet0/0  - LAN A  encrypted
ip address 10.3.1.1 255.255.255.192

interface FastEthernet0/1  - LAN B not encrypted
ip address 10.3.2.1 255.255.255.240

interface FastEthernet1/0  - To NCL PE
crypto map RD-EXEC-CRYPTOMAP  - Same as above
```

Testing Evidence

Test 3.1: VRF-EXEC-RD Routing Table

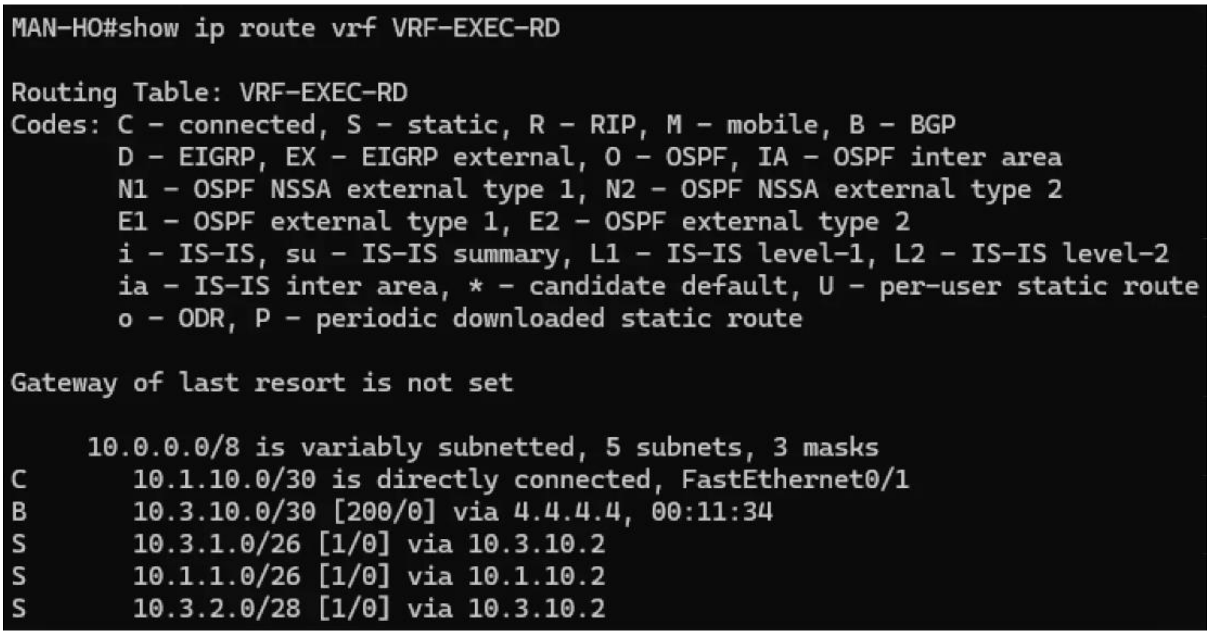

```
MAN-HO#show ip route vrf VRF-EXEC-RD

Routing Table: VRF-EXEC-RD
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
C       10.1.10.0/30 is directly connected, FastEthernet0/1
B       10.3.10.0/30 [200/0] via 4.4.4.4, 00:11:34
S       10.3.1.0/26 [1/0] via 10.3.10.2
S       10.1.1.0/26 [1/0] via 10.1.10.2
S       10.3.2.0/28 [1/0] via 10.3.10.2
```

Explanation: Shows routes to EXEC LAN (10.1.1.0/26), R-and-D LAN A (10.3.1.0/26), and LAN B (10.3.2.0/28).

Test 3.2: BGP VPNv4 for VRF-EXEC-RD

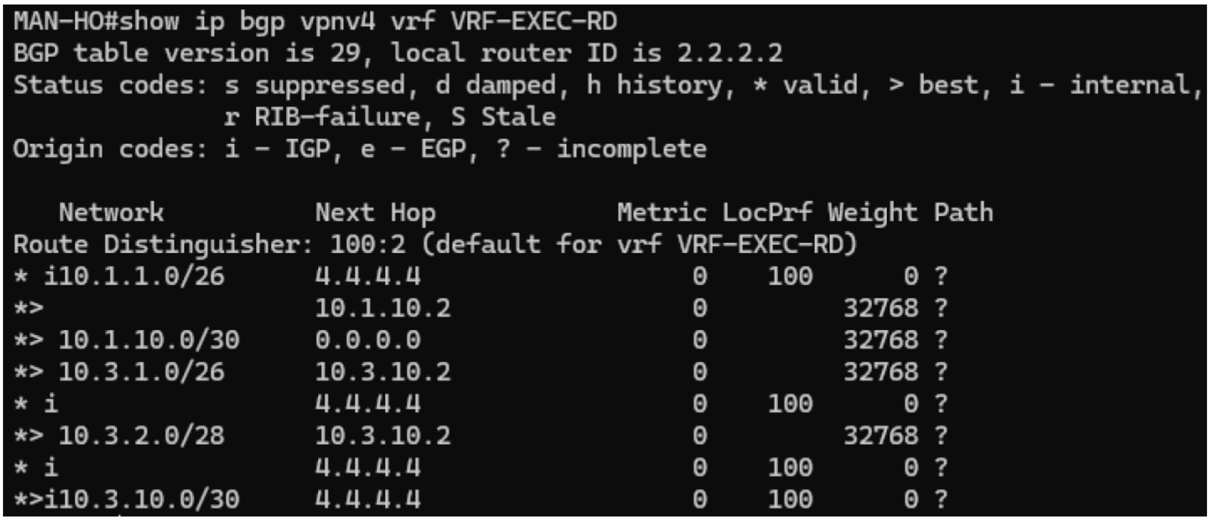

```
MAN-HO#show ip bgp vpnv4 vrf VRF-EXEC-RD
BGP table version is 29, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:2 (default for vrf VRF-EXEC-RD)
* i10.1.1.0/26      4.4.4.4                  0    100      0 ?
*>                  10.1.10.2                0         32768 ?
*> 10.1.10.0/30     0.0.0.0                  0         32768 ?
*> 10.3.1.0/26      10.3.10.2                0         32768 ?
* i                 4.4.4.4                  0    100      0 ?
*> 10.3.2.0/28      10.3.10.2                0         32768 ?
* i                 4.4.4.4                  0    100      0 ?
*>i10.3.10.0/30     4.4.4.4                  0    100      0 ?
```

Explanation: Shows routes with RD 100:2 being exchanged between MAN-HO and NCL.

Test 3.3: IKE Security Association

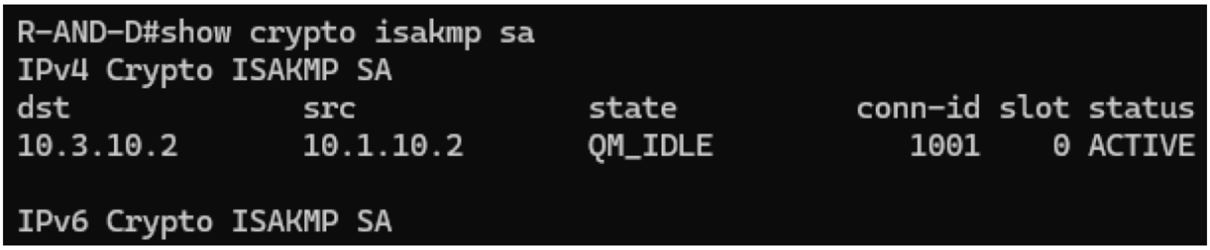

```
R-AND-D#show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst             src             state          conn-id slot status
10.3.10.2       10.1.10.2       QM_IDLE           1001    0 ACTIVE

IPv6 Crypto ISAKMP SA
```

Explanation: Shows IKE SA between R-AND-D and EXEC in QM_IDLE state.

Test 3.4: IPSec Security Association

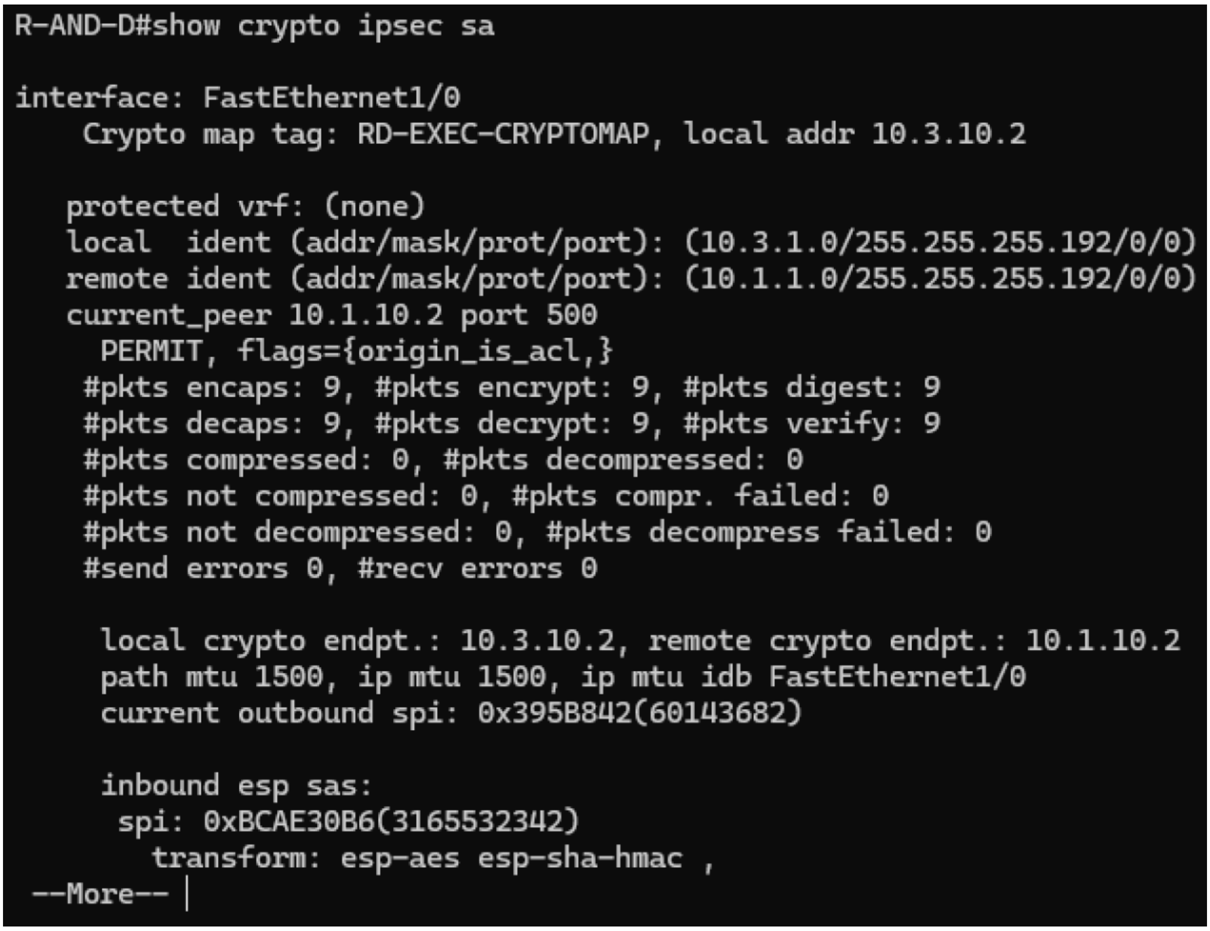

```
R-AND-D#show crypto ipsec sa

interface: FastEthernet1/0
    Crypto map tag: RD-EXEC-CRYPTOMAP, local addr 10.3.10.2

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (10.3.1.0/255.255.255.192/0/0)
   remote ident (addr/mask/prot/port): (10.1.1.0/255.255.255.192/0/0)
   current_peer 10.1.10.2 port 500
     PERMIT, flags={origin_is_acl,}
    #pkts encaps: 9, #pkts encrypt: 9, #pkts digest: 9
    #pkts decaps: 9, #pkts decrypt: 9, #pkts verify: 9
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 10.3.10.2, remote crypto endpt.: 10.1.10.2
     path mtu 1500, ip mtu 1500, ip mtu idb FastEthernet1/0
     current outbound spi: 0x395B842(60143682)

     inbound esp sas:
      spi: 0xBCAE30B6(3165532342)
        transform: esp-aes esp-sha-hmac ,
 --More-- |
```

Explanation: Shows IPSec SA with ACL matching only 10.3.1.0/26 (LAN A). Packet counters shows encrypted traffic.

Test 3.5: Ping to LAN A (Encrypted)

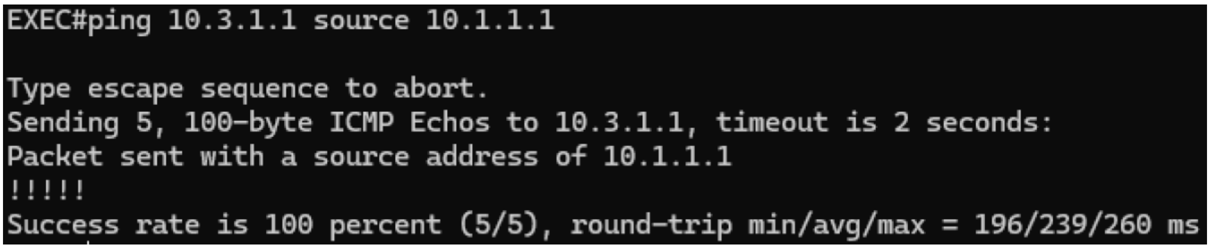

```
EXEC#ping 10.3.1.1 source 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.1.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 196/239/260 ms
```

Explanation: 100% success shows this traffic is encrypted (matches crypto ACL).

Test 3.6: Ping to LAN B (Unencrypted)

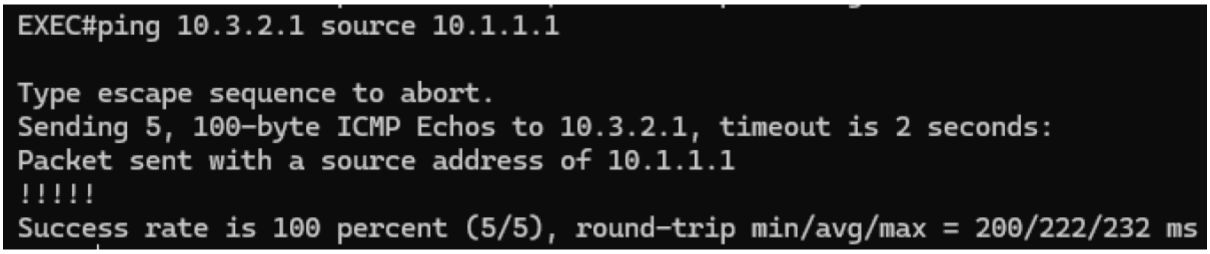

```
EXEC#ping 10.3.2.1 source 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.2.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 200/222/232 ms
```

Explanation: 100% success. This traffic is NOT encrypted (LAN B excluded from ACL).

Test 3.7: Traceroute to LAN A

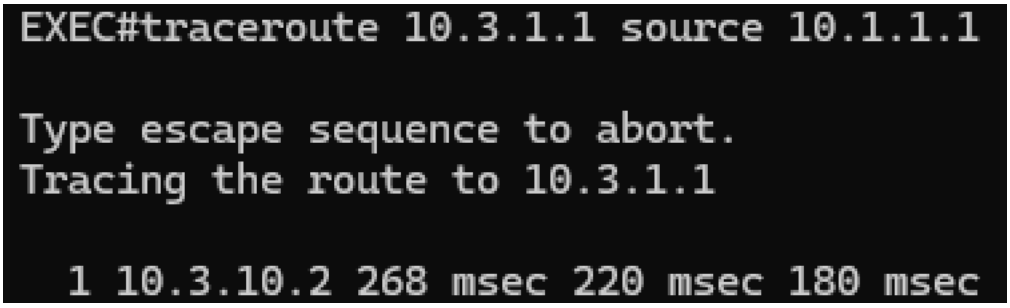

```
EXEC#traceroute 10.3.1.1 source 10.1.1.1

Type escape sequence to abort.
Tracing the route to 10.3.1.1

  1 10.3.10.2 268 msec 220 msec 180 msec
```

Explanation: Shows traffic path to R AND D LAN A via MPLS VPN. The first hop (10.3.10.2) is the R AND D router. This traffic is IPSec encrypted as it matches the crypto ACL for LAN A.

2.4 Requirement 4: SQL Server Access Control at LOGISTICS

Router Configurations

LOGISTICS Router - Extended ACL:

ip access-list extended SQL-SERVER-ACCESS *- Named extended ACL*
permit tcp 10.1.2.0 0.0.0.127 host 10.3.4.5 eq 1433 *- Allows HR subnet to connect to SQL server on port 1433*
permit tcp host 10.3.4.5 eq 1433 10.1.2.0 0.0.0.127 established *- Allows return traffic to HR*
deny ip any host 10.3.4.5 log *- Denies everything else and logs the attempts*
permit ip any any *- Allows all non SQL traffic to pass normally*

interface FastEthernet0/1 *- Interface to NCL PE*
ip access-group SQL-SERVER-ACCESS in *- Filters traffic entering from PE router*

ACL Logic:

Line 1 permits HR (10.1.2.0/25) TCP connections to SQL server (10.3.4.5) port 1433. Line 2 allows return traffic. Line 3 denies and logs all other access to SQL server. Line 4 permits everything else so normal traffic is not blocked.

Testing Evidence

Test 4.1: Running Configuration - ACL

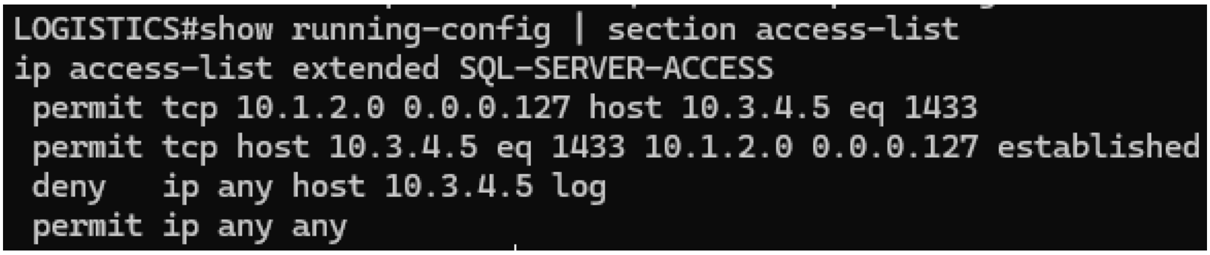

```
LOGISTICS#show running-config | section access-list
ip access-list extended SQL-SERVER-ACCESS
 permit tcp 10.1.2.0 0.0.0.127 host 10.3.4.5 eq 1433
 permit tcp host 10.3.4.5 eq 1433 10.1.2.0 0.0.0.127 established
 deny   ip any host 10.3.4.5 log
 permit ip any any
```

Explanation: Shows the SQL-SERVER-ACCESS ACL configuration with all four entries.

Test 4.2: Access List Counters (Before Tests)

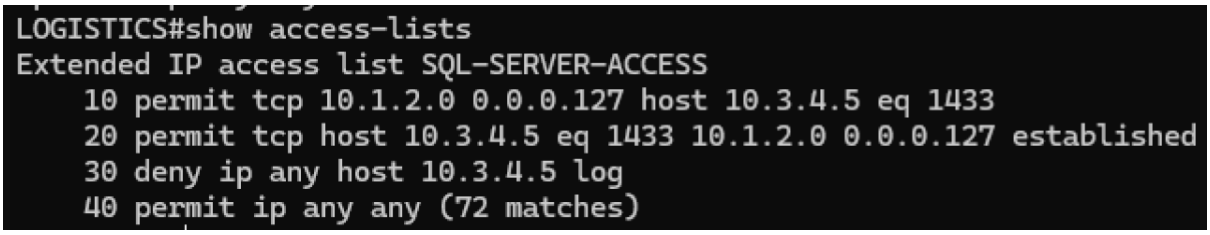

```
LOGISTICS#show access-lists
Extended IP access list SQL-SERVER-ACCESS
    10 permit tcp 10.1.2.0 0.0.0.127 host 10.3.4.5 eq 1433
    20 permit tcp host 10.3.4.5 eq 1433 10.1.2.0 0.0.0.127 established
    30 deny ip any host 10.3.4.5 log
    40 permit ip any any (72 matches)
```

Explanation: Shows current match counters for each ACL line. Note the values before testing.

Test 4.3: HR Access Attempt to SQL Port 1433

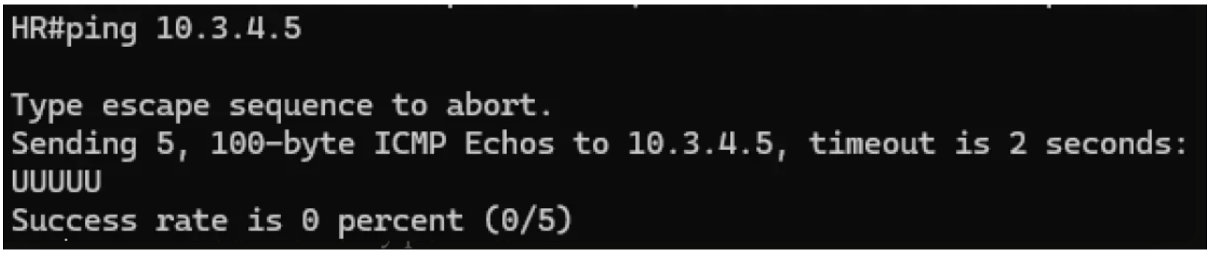

```
HR#ping 10.3.4.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.4.5, timeout is 2 seconds:
UUUUU
Success rate is 0 percent (0/5)
```

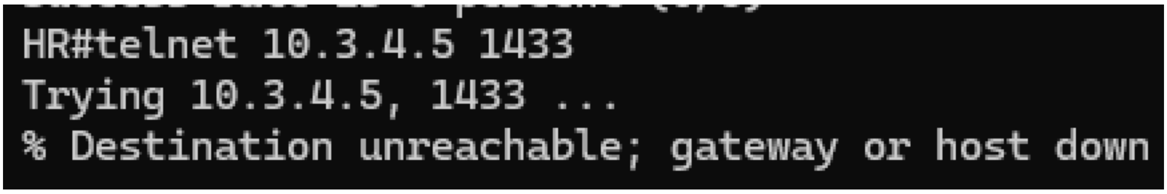

```
HR#telnet 10.3.4.5 1433
Trying 10.3.4.5, 1433 ...
% Destination unreachable; gateway or host down
```

Explanation: The ping (ICMP) fails because the ACL only permits TCP port 1433, not ICMP. The telnet to port 1433 shows "Destination unreachable" because no actual SQL server exists in the lab environment, but the traffic is permitted by ACL line 10. The ACL counter will confirm this traffic was allowed.

Test 4.4: Access List Counters (After HR Test)

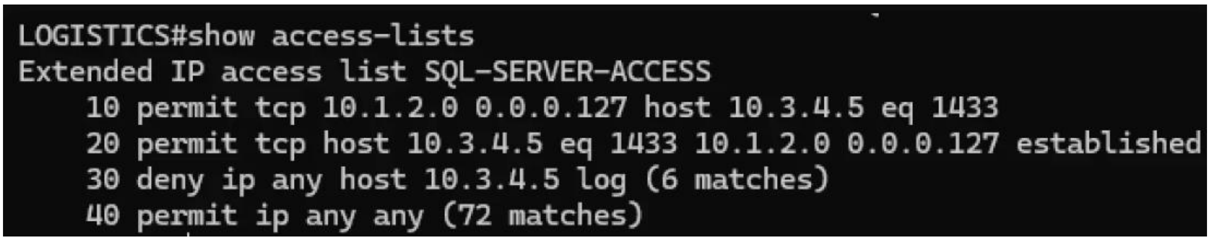

```
LOGISTICS#show access-lists
Extended IP access list SQL-SERVER-ACCESS
    10 permit tcp 10.1.2.0 0.0.0.127 host 10.3.4.5 eq 1433
    20 permit tcp host 10.3.4.5 eq 1433 10.1.2.0 0.0.0.127 established
    30 deny ip any host 10.3.4.5 log (6 matches)
    40 permit ip any any (72 matches)
```

Explanation: Line 30 shows 6 matches. 5 from the blocked ICMP ping and 1 from the telnet attempt. The telnet was also denied because the router's source IP is not within the HR LAN subnet (10.1.2.0/25). This proves the ACL correctly restricts access to only hosts on the HR LAN.

Test 4.5: EXEC Access Attempt to SQL Port 1433

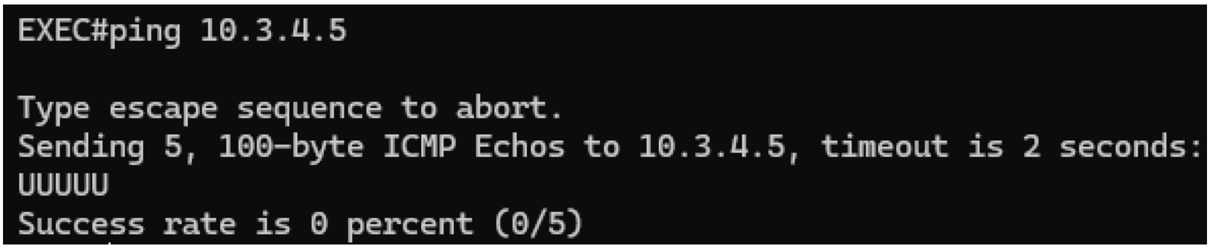

```
EXEC#ping 10.3.4.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.4.5, timeout is 2 seconds:
UUUUU
Success rate is 0 percent (0/5)
```

Explanation: Connection denied by ACL

Test 4.6: Access List Counters (After EXEC Test)

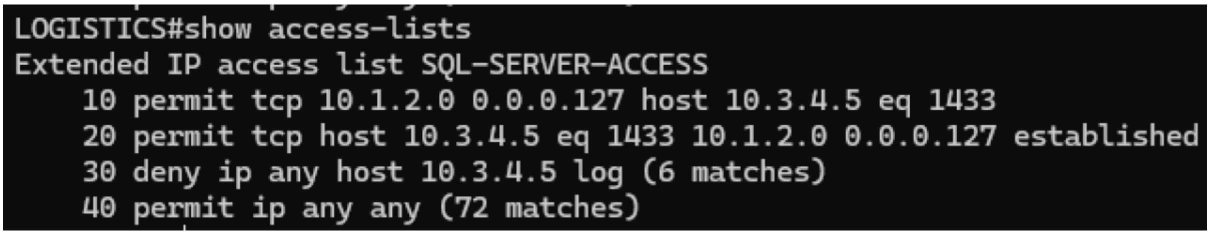

```
LOGISTICS#show access-lists
Extended IP access list SQL-SERVER-ACCESS
    10 permit tcp 10.1.2.0 0.0.0.127 host 10.3.4.5 eq 1433
    20 permit tcp host 10.3.4.5 eq 1433 10.1.2.0 0.0.0.127 established
    30 deny ip any host 10.3.4.5 log (6 matches)
    40 permit ip any any (72 matches)
```

Explanation: Counter unchanged because EXEC is in VRF-EXEC-RD, which has no route to the SQL server in VRF-MESH. EXEC is blocked by VRF separation, not by the ACL

Test 4.7: Ping from PROD-1 to SQL Server (Blocked)

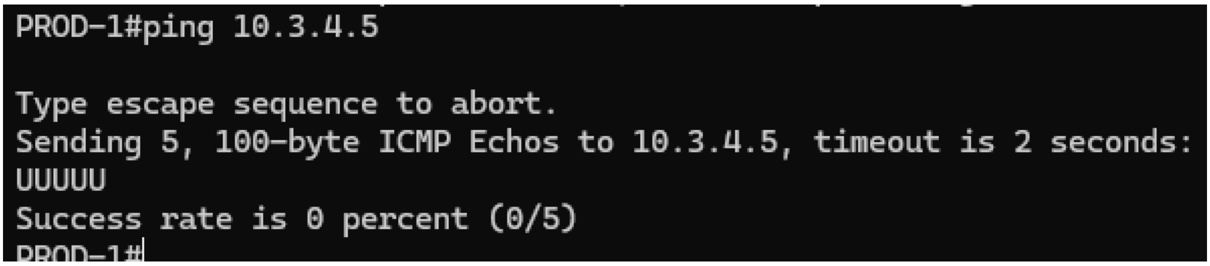

```
PROD-1#ping 10.3.4.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.4.5, timeout is 2 seconds:
UUUUU
Success rate is 0 percent (0/5)
PROD-1#
```

Explanation: PROD-1 is in VRF-MESH so traffic reaches LOGISTICS but is blocked by ACL in line 30.

Test 4.8: Access List Counters (After PROD-1 Test)

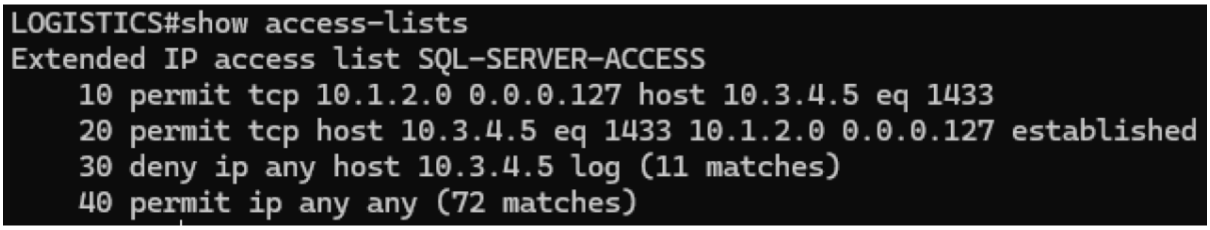

```
LOGISTICS#show access-lists
Extended IP access list SQL-SERVER-ACCESS
    10 permit tcp 10.1.2.0 0.0.0.127 host 10.3.4.5 eq 1433
    20 permit tcp host 10.3.4.5 eq 1433 10.1.2.0 0.0.0.127 established
    30 deny ip any host 10.3.4.5 log (11 matches)
    40 permit ip any any (72 matches)
```

Explanation: Line 30 counter increased, proving PROD-1 traffic reached LOGISTICS and was blocked by the ACL.

Test 4.9: Traceroute from HR to SQL Server

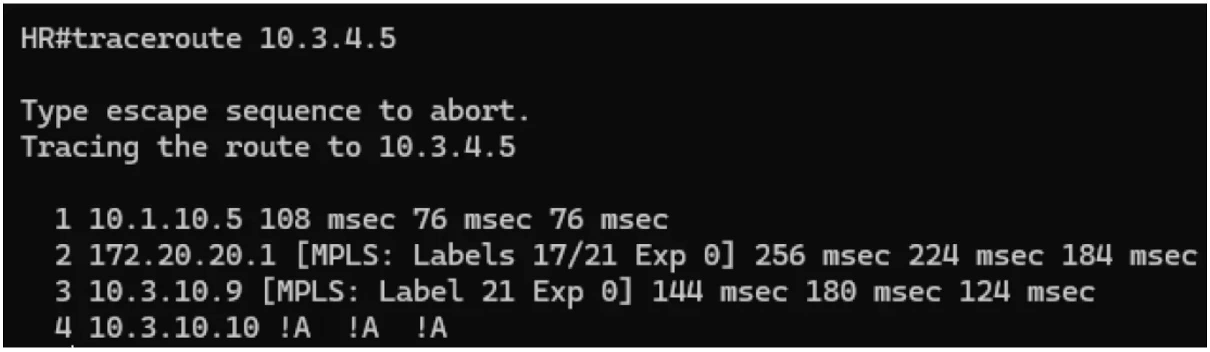

```
HR#traceroute 10.3.4.5

Type escape sequence to abort.
Tracing the route to 10.3.4.5

  1 10.1.10.5 108 msec 76 msec 76 msec
  2 172.20.20.1 [MPLS: Labels 17/21 Exp 0] 256 msec 224 msec 184 msec
  3 10.3.10.9 [MPLS: Label 21 Exp 0] 144 msec 180 msec 124 msec
  4 10.3.10.10 !A  !A  !A
```

Explanation: Shows path through MPLS backbone to LOGISTICS. The !A (administratively prohibited) at hop 4 confirms the ACL is blocking traffic at the destination while the route exists.

2.5 Requirement 5: PROD-2 Isolation from EXEC

Router Configurations

NCL Router (PE) - VRF Configuration:

```
ip vrf VRF-EXEC-RD  - EXEC is in this VRF
rd 100:2 - Same as above
route-target export 100:2 - Same as above
route-target import 100:2 - Same as above

ip vrf VRF-MESH  - LOGISTICS is in this VRF
rd 100:1 - Same as above
route-target export 100:1 - Same as above
route-target import 100:1 - Same as above

interface FastEthernet0/1  - To R-AND-D (VRF-EXEC-RD)
ip vrf forwarding VRF-EXEC-RD - Same as above
ip address 10.3.10.1 255.255.255.252 - Same as above

interface FastEthernet1/0  - To PROD-2 (Global table - no VRF)
ip address 10.3.10.5 255.255.255.252  - NOT in any VRF

interface FastEthernet2/0  - To LOGISTICS (VRF-MESH)
ip vrf forwarding VRF-MESH - Same as above
ip address 10.3.10.9 255.255.255.252 - Same as above
```

PROD-2 connects to NCL's global routing table (no VRF assigned to FastEthernet1/0). EXEC is in VRF-EXEC-RD with RD 100:2. Since these are completely separate routing domains, no routes are exchanged between them. PROD-2 has no route to reach EXEC, and vice versa.

Testing Evidence

Test 5.1: Running Configuration - VRF Interfaces on NCL

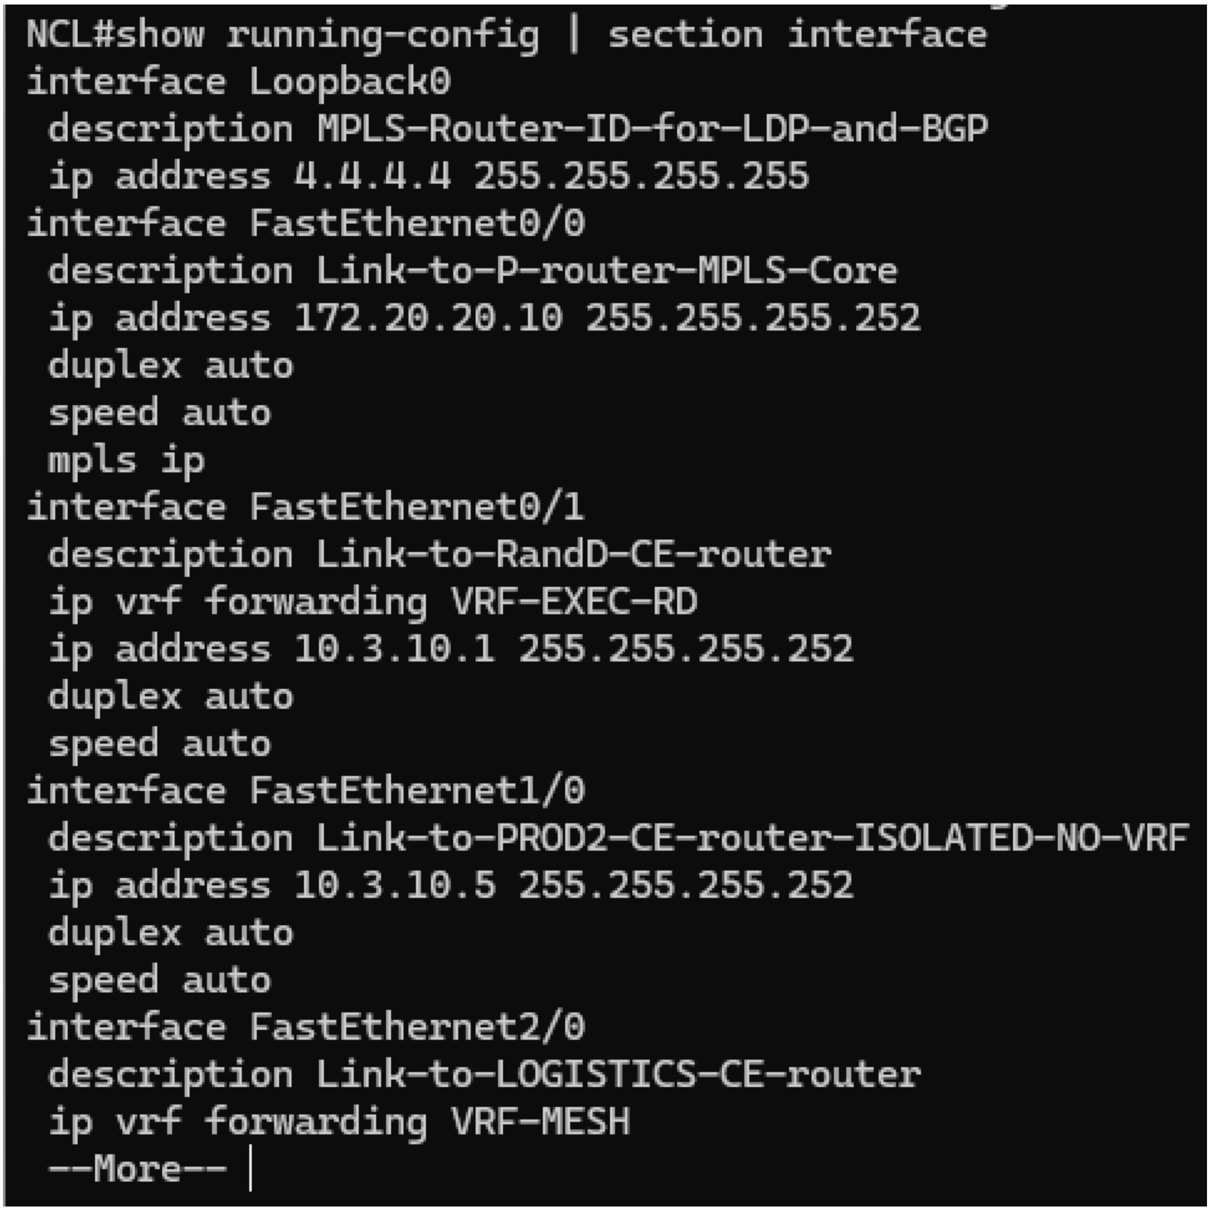

```
NCL#show running-config | section interface
interface Loopback0
 description MPLS-Router-ID-for-LDP-and-BGP
 ip address 4.4.4.4 255.255.255.255
interface FastEthernet0/0
 description Link-to-P-router-MPLS-Core
 ip address 172.20.20.10 255.255.255.252
 duplex auto
 speed auto
 mpls ip
interface FastEthernet0/1
 description Link-to-RandD-CE-router
 ip vrf forwarding VRF-EXEC-RD
 ip address 10.3.10.1 255.255.255.252
 duplex auto
 speed auto
interface FastEthernet1/0
 description Link-to-PROD2-CE-router-ISOLATED-NO-VRF
 ip address 10.3.10.5 255.255.255.252
 duplex auto
 speed auto
interface FastEthernet2/0
 description Link-to-LOGISTICS-CE-router
 ip vrf forwarding VRF-MESH
 --More-- |
```

Explanation: Shows FastEthernet1/0 (to PROD-2) has no 'ip vrf forwarding' command, while other interfaces have VRF assignments.

Test 5.2: PROD-2 Routing Table

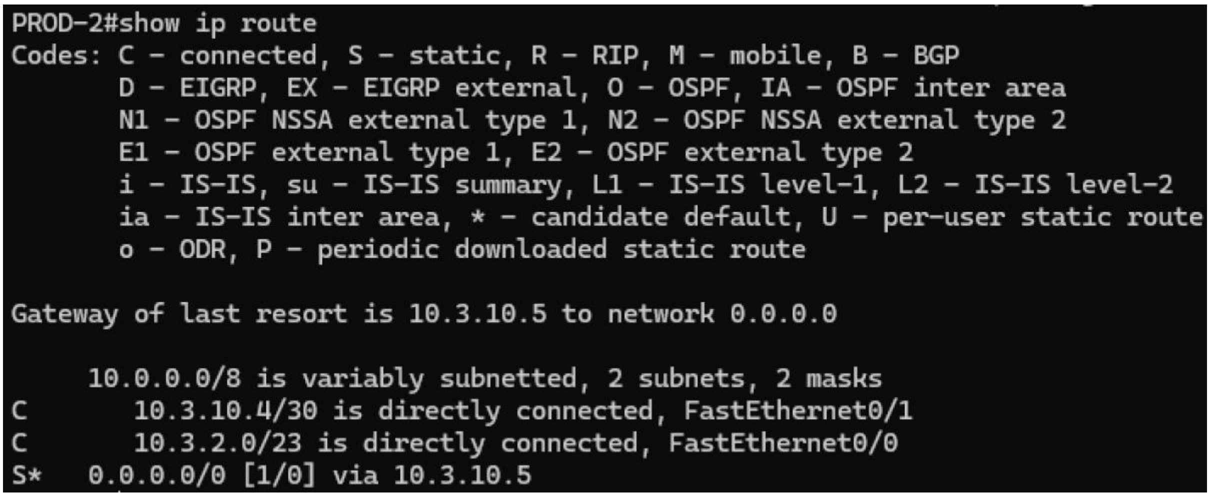

```
PROD-2#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 10.3.10.5 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C       10.3.10.4/30 is directly connected, FastEthernet0/1
C       10.3.2.0/23 is directly connected, FastEthernet0/0
S*   0.0.0.0/0 [1/0] via 10.3.10.5
```

Explanation: Shows PROD-2's routes. There is NO route to 10.1.1.0/26 (EXEC LAN) because it exists in a different routing domain.

Test 5.3: Ping from PROD-2 to EXEC

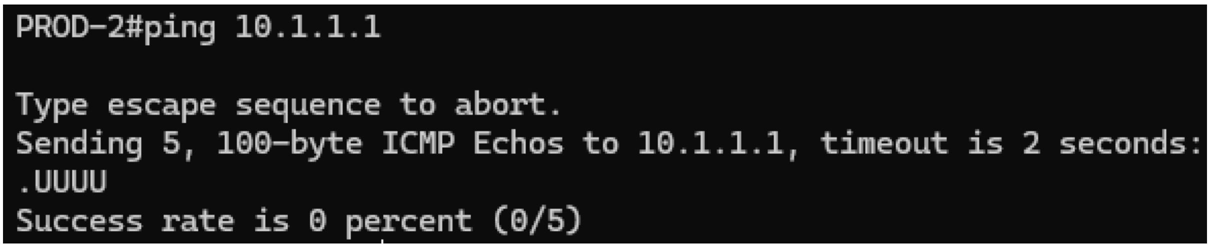

```
PROD-2#ping 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
.UUUU
Success rate is 0 percent (0/5)
```

Explanation: Fails with UUUUU (0% success). Destination unreachable, no route exists.

Test 5.4: Traceroute from PROD-2 to EXEC

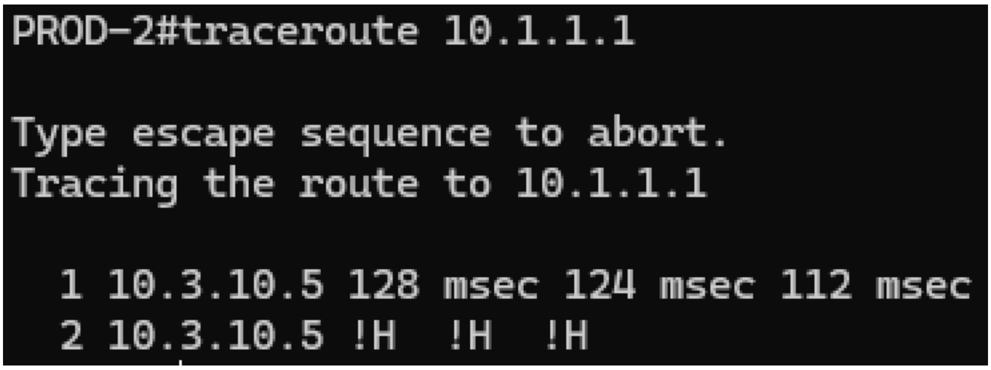

```
PROD-2#traceroute 10.1.1.1

Type escape sequence to abort.
Tracing the route to 10.1.1.1

  1 10.3.10.5 128 msec 124 msec 112 msec
  2 10.3.10.5 !H  !H  !H
```

Explanation: Shows packets reach gateway (10.3.10.5) then fail with !H (host unreachable). No further paths exist.

Test 5.5: Ping from EXEC to PROD-2

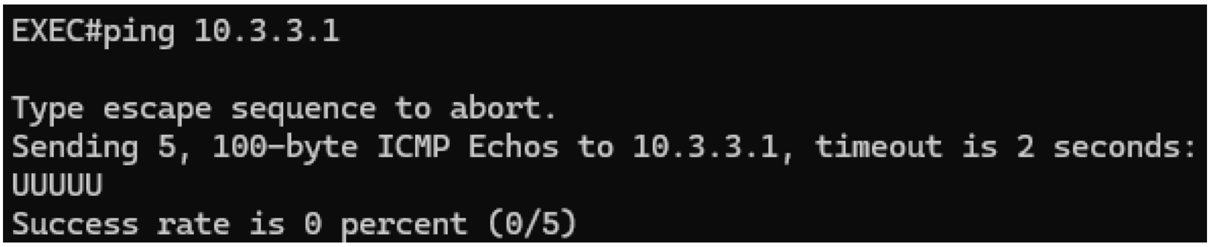

```
EXEC#ping 10.3.3.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.3.1, timeout is 2 seconds:
UUUUU
Success rate is 0 percent (0/5)
```

Explanation: Fails (0% success). Confirms isolation works in both directions.

Test 5.6: PROD-2 Connectivity to Gateway

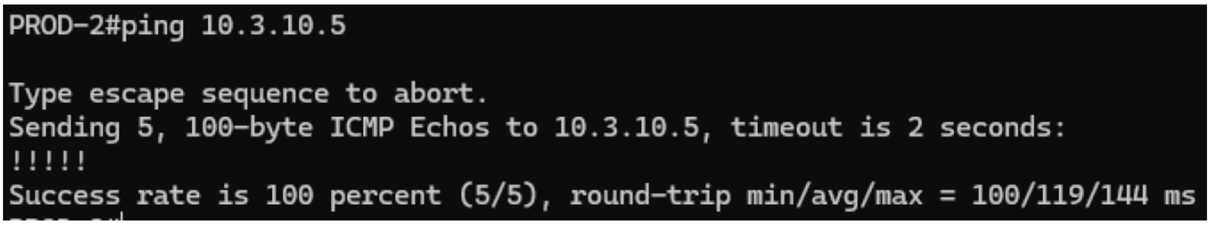

```
PROD-2#ping 10.3.10.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.10.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/119/144 ms
```

Explanation: Succeeds 100% which proves PROD-2 has basic connectivity. Isolation is VRF-based.

2.6 Requirement 6: DHCP Services Across All LANs

Router Configurations

EXEC Router:

```
ip dhcp excluded-address 10.1.1.1 10.1.1.9  - Reserve infrastructure Ips, excludes first 9 addresses from DHCP pool
ip dhcp pool EXEC-LAN – Creates DHCP pool for EXEC
network 10.1.1.0 255.255.255.192  - /26 = 62 addresses
default-router 10.1.1.1 – Gateway address given to clients
dns-server 8.8.8.8 8.8.4.4 – Google DNS server for name resolutions
lease 7 – Ip addresses valid for 7 days
```

HR Router:

```
ip dhcp excluded-address 10.1.2.1 10.1.2.9 - Same as above
ip dhcp pool HR-LAN - Same as above
network 10.1.2.0 255.255.255.128  - /25 = 126 addresses
default-router 10.1.2.1 - Same as above
dns-server 8.8.8.8 8.8.4.4 - Same as above
lease 7 - Same as above
```

PROD-1 Router:

```
ip dhcp excluded-address 10.2.1.1 10.2.1.9 - Same as above
ip dhcp pool PROD1-LAN - Same as above
network 10.2.1.0 255.255.255.128  - /25 = 126 addresses
default-router 10.2.1.1 - Same as above
dns-server 8.8.8.8 8.8.4.4 - Same as above
lease 7 - Same as above
```

R-AND-D Router (Two LANs):

```
ip dhcp excluded-address 10.3.1.1 10.3.1.9  - LAN A exclusions
ip dhcp excluded-address 10.3.2.1  - LAN B gateway
ip dhcp pool RD-LAN-A - Same as above
network 10.3.1.0 255.255.255.192  - /26 = 62 addresses
default-router 10.3.1.1 - Same as above
dns-server 8.8.8.8 8.8.4.4 - Same as above
lease 7 - Same as above
ip dhcp pool RD-LAN-B - Same as above
network 10.3.2.0 255.255.255.240  - /28 = 14 addresses
default-router 10.3.2.1 - Same as above
dns-server 8.8.8.8 8.8.4.4 - Same as above
lease 7 - Same as above
```

PROD-2 Router:

```
ip dhcp excluded-address 10.3.3.1 10.3.3.9 - Same as above
ip dhcp pool PROD2-LAN - Same as above
network 10.3.2.0 255.255.254.0  - /23 = 510 addresses
default-router 10.3.3.1 - Same as above
dns-server 8.8.8.8 8.8.4.4 - Same as above
lease 7 - Same as above
```

LOGISTICS Router:

ip dhcp excluded-address 10.3.4.1 10.3.4.9 *- Infrastructure, reserves addresses for network devices*
ip dhcp excluded-address 10.3.4.5 *- SQL server reserved, out of DHCP pool*
ip dhcp pool LOGISTICS-LAN *- Same as above*
network 10.3.4.0 255.255.255.128 *- /25 = 126 addresses*
default-router 10.3.4.1 *- Same as above*
dns-server 8.8.8.8 8.8.4.4 *- Same as above*
lease 7 *- Same as above*

Testing Evidence

Test 6.1: DHCP Pool on EXEC

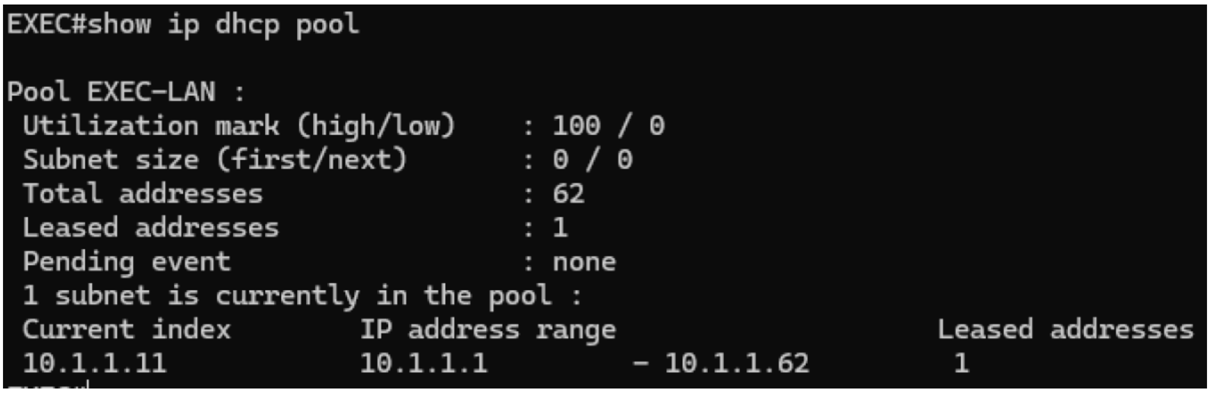

```
EXEC#show ip dhcp pool

Pool EXEC-LAN :
 Utilization mark (high/low)    : 100 / 0
 Subnet size (first/next)       : 0 / 0
 Total addresses                : 62
 Leased addresses               : 1
 Pending event                  : none
 1 subnet is currently in the pool :
 Current index        IP address range                    Leased addresses
 10.1.1.11            10.1.1.1           - 10.1.1.62       1
```

Explanation: Shows EXEC-LAN pool with network 10.1.1.0/26, 62 total addresses. Leased addresses shows 1, confirming DHCP is actively assigning IPs to clients.

Test 6.2: DHCP Binding on EXEC

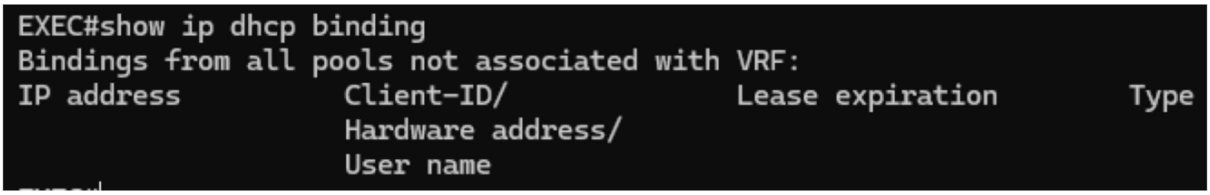

```
EXEC#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
```

Explanation: Shows active lease for 10.1.1.10 assigned to a client. This confirms DHCP server is functioning and distributing addresses from the configured pool.

Test 6.3: DHCP Pool on HR

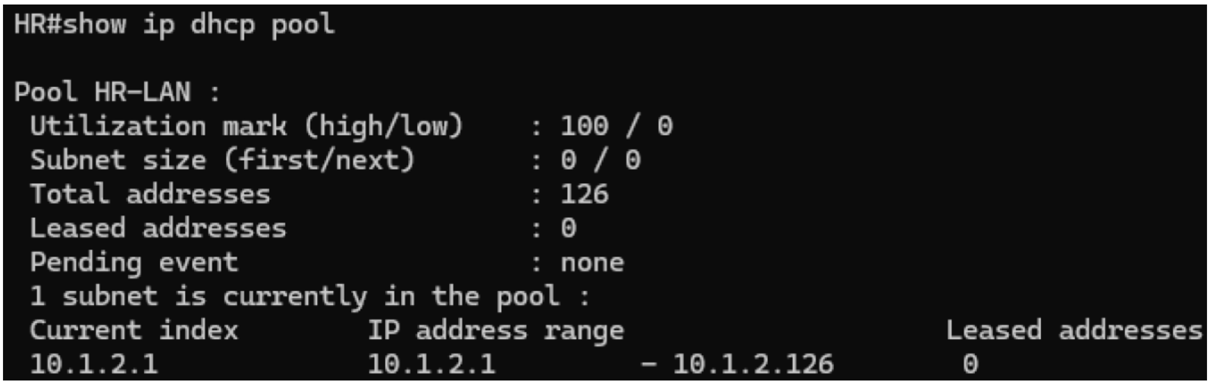

```
HR#show ip dhcp pool

Pool HR-LAN :
 Utilization mark (high/low)    : 100 / 0
 Subnet size (first/next)       : 0 / 0
 Total addresses                : 126
 Leased addresses               : 0
 Pending event                  : none
 1 subnet is currently in the pool :
 Current index        IP address range                    Leased addresses
 10.1.2.1             10.1.2.1           - 10.1.2.126      0
```

Explanation: Shows HR-LAN pool with network 10.1.2.0/25, 126 total addresses.

Test 6.4: DHCP Pool on PROD-1

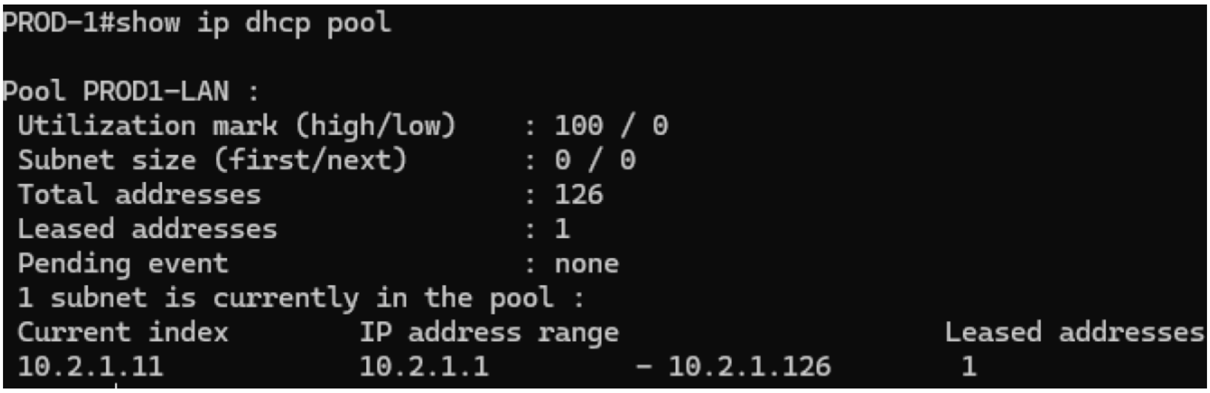

```
PROD-1#show ip dhcp pool

Pool PROD1-LAN :
 Utilization mark (high/low)    : 100 / 0
 Subnet size (first/next)       : 0 / 0
 Total addresses                : 126
 Leased addresses               : 1
 Pending event                  : none
 1 subnet is currently in the pool :
 Current index        IP address range                    Leased addresses
 10.2.1.11            10.2.1.1         - 10.2.1.126        1
```

Explanation: Shows PROD1-LAN pool with network 10.2.1.0/25, 126 total addresses.

Test 6.5: DHCP Binding on PROD-1

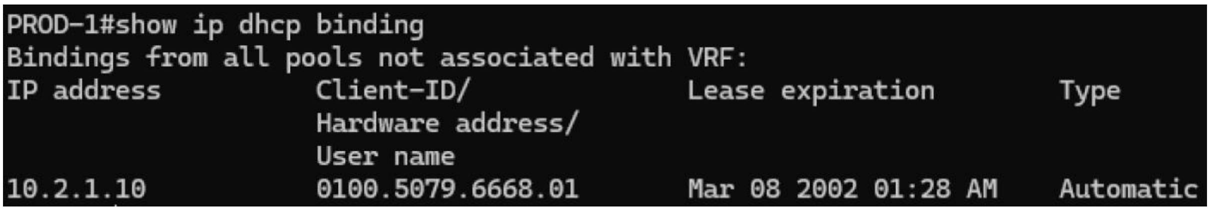

```
PROD-1#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
10.2.1.10           0100.5079.6668.01       Mar 08 2002 01:28 AM    Automatic
```

Explanation: Shows active lease for 10.2.1.10 confirming DHCP server is distributing addresses at the LDS site.

Test 6.6: DHCP Pools on R-AND-D (Two Pools)

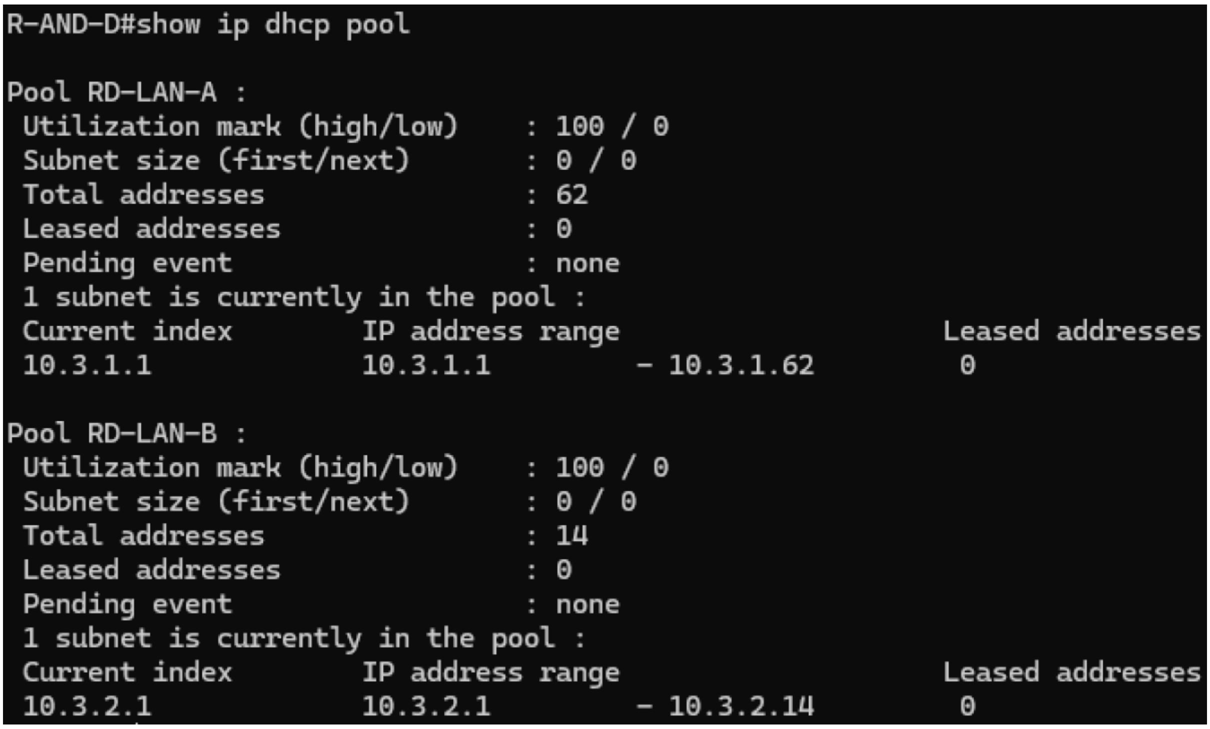

```
R-AND-D#show ip dhcp pool

Pool RD-LAN-A :
 Utilization mark (high/low)    : 100 / 0
 Subnet size (first/next)       : 0 / 0
 Total addresses                : 62
 Leased addresses               : 0
 Pending event                  : none
 1 subnet is currently in the pool :
 Current index        IP address range                    Leased addresses
 10.3.1.1             10.3.1.1         - 10.3.1.62         0

Pool RD-LAN-B :
 Utilization mark (high/low)    : 100 / 0
 Subnet size (first/next)       : 0 / 0
 Total addresses                : 14
 Leased addresses               : 0
 Pending event                  : none
 1 subnet is currently in the pool :
 Current index        IP address range                    Leased addresses
 10.3.2.1             10.3.2.1         - 10.3.2.14         0
```

Explanation: Shows BOTH pools: RD-LAN-A (10.3.1.0/26, 62 addr) and RD-LAN-B (10.3.2.0/28, 14 addr).

Test 6.7: DHCP Pool on PROD-2

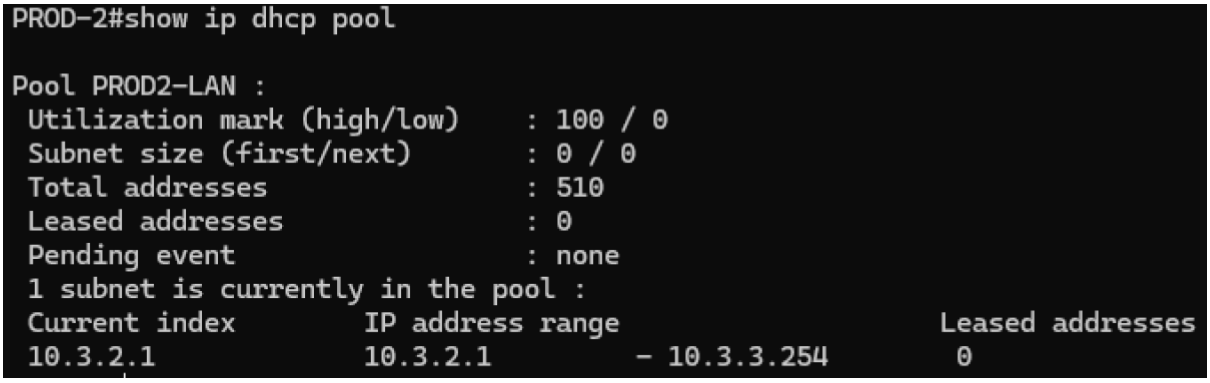

```
PROD-2#show ip dhcp pool

Pool PROD2-LAN :
 Utilization mark (high/low)    : 100 / 0
 Subnet size (first/next)       : 0 / 0
 Total addresses                : 510
 Leased addresses               : 0
 Pending event                  : none
 1 subnet is currently in the pool :
 Current index        IP address range                    Leased addresses
 10.3.2.1             10.3.2.1         - 10.3.3.254        0
```

Explanation: Shows PROD2-LAN pool with network 10.3.2.0/23, 510 total addresses.

Test 6.8: DHCP Pool on LOGISTICS

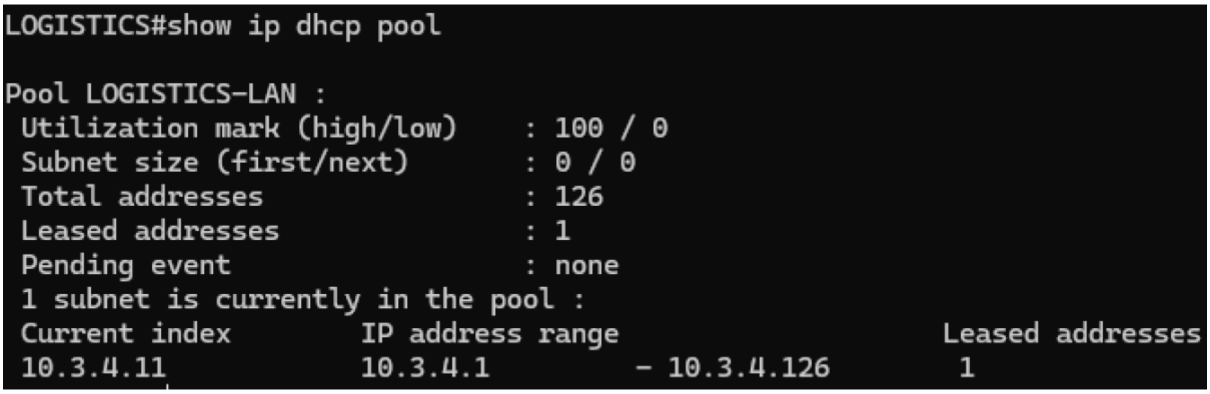

```
LOGISTICS#show ip dhcp pool

Pool LOGISTICS-LAN :
 Utilization mark (high/low)    : 100 / 0
 Subnet size (first/next)       : 0 / 0
 Total addresses                : 126
 Leased addresses               : 1
 Pending event                  : none
 1 subnet is currently in the pool :
 Current index        IP address range                    Leased addresses
 10.3.4.11            10.3.4.1         - 10.3.4.126        1
```

Explanation: Shows LOGISTICS-LAN pool with network 10.3.4.0/25, 126 total addresses. Leased addresses shows 1, confirming active DHCP at NCL site.

Test 6.9: DHCP Binding on LOGISTICS

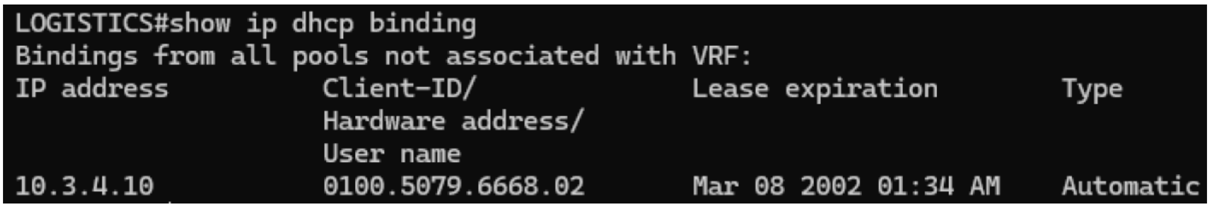

```
LOGISTICS#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
10.3.4.10           0100.5079.6668.02       Mar 08 2002 01:34 AM    Automatic
```

Explanation: Shows active lease for 10.3.4.10 confirming DHCP server is distributing addresses at the NCL site.

Test 6.10: Gateway Connectivity Test

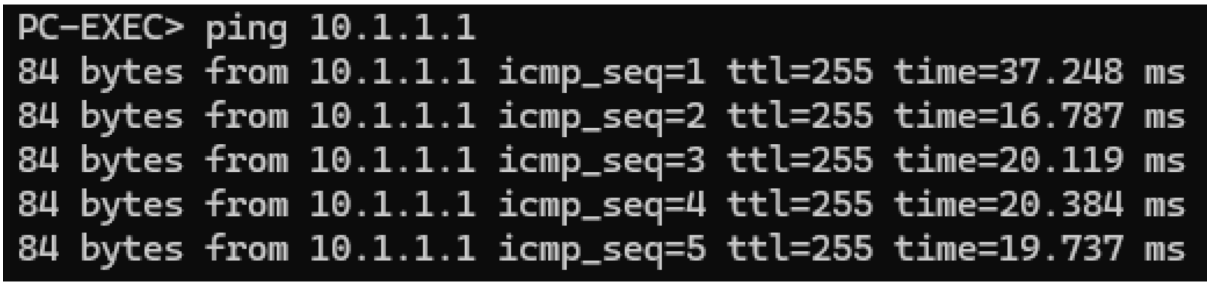

```
PC-EXEC> ping 10.1.1.1
84 bytes from 10.1.1.1 icmp_seq=1 ttl=255 time=37.248 ms
84 bytes from 10.1.1.1 icmp_seq=2 ttl=255 time=16.787 ms
84 bytes from 10.1.1.1 icmp_seq=3 ttl=255 time=20.119 ms
84 bytes from 10.1.1.1 icmp_seq=4 ttl=255 time=20.384 ms
84 bytes from 10.1.1.1 icmp_seq=5 ttl=255 time=19.737 ms
```

Explanation: VPC successfully pings gateway (10.1.1.1), confirms DHCP clients receive correct gateway and have connectivity.

DHCP Summary Table

Router	Pool Name	Network	Mask	Total	Gateway
EXEC	EXEC-LAN	10.1.1.0	/26	62	10.1.1.1
HR	HR-LAN	10.1.2.0	/25	126	10.1.2.1
PROD-1	PROD1-LAN	10.2.1.0	/25	126	10.2.1.1
R-AND-D	RD-LAN-A	10.3.1.0	/26	62	10.3.1.1
R-AND-D	RD-LAN-B	10.3.2.0	/28	14	10.3.2.1
PROD-2	PROD2-LAN	10.3.2.0	/23	510	10.3.3.1
LOGISTICS	LOGISTICS-LAN	10.3.4.0	/25	126	10.3.4.1

All 7 DHCP pools are configured correctly with subnet sizes matching department requirements, with the distributed architecture ensuring site independence.

NOTE: DHCP binding tests were conducted using VPCs at three representative sites: EXEC (MAN_HO), PROD-1 (LDS), and LOGISTICS (NCL). This demonstrates DHCP functionality across all three geographic locations without redundant testing. Pool configuration screenshots for all routers confirm DHCP is properly configured network-wide.

Section 3: Strategy for Implementation Plan and Troubleshooting

3.1 Implementation Plan

Phase 1: MPLS Core - P Router

The P router forms the backbone for all MPLS services and is configured first. This involves assigning the loopback address, configuring OSPF Area 0 on all interfaces, and enabling MPLS IP on each PE-facing interface. The loopback provides a stable router-id unaffected by physical interface failures, which is essential for both OSPF and LDP stability.

Phase 2: MPLS Core - PE Routers

Each PE router receives loopback configuration and OSPF Area 0 membership. MPLS IP is enabled on provider-facing interfaces using the specified addressing: MAN-HO at 172.20.20.2/30, LDS at 172.20.20.6/30, and NCL at 172.20.20.10/30. LDP is enabled using *mpls ldp router-id Loopback0 force* on all MPLS routers.

Phase 3: MPLS Core Verification

Before proceeding, the MPLS backbone is verified. OSPF adjacencies are confirmed using *show ip ospf neighbor,* ensuring all PE-P relationships show FULL state. LDP sessions are verified using *show mpls ldp neighbor*, confirming three peers in Oper state on the P router. The command *show mpls forwarding-table* confirms labels are assigned to all loopbacks. This verification step prevents failures in later phases.

Phase 4: BGP VPNv4

With MPLS operational, iBGP sessions are established between PE routers using loopback addresses as update sources for stability. All PEs are configured with AS 65000. The VPNv4 address family is activated with *send-community both* to ensure route target communities propagate correctly. Verification using *show ip bgp summary* confirms all peers reach established state.

Phase 5: VRF Definitions

VRF instances are created on PE routers before any customer-facing interface configuration. VRF-MESH (RD 100:1, RT 100:1 import/export) is defined on PE routers for the HR, PROD-1, and LOGISTICS sites. VRF-EXEC-RD (RD 100:2, RT 100:2 import/export) is defined on MAN-HO and NCL only for EXEC and R-and-D connectivity. Creating VRFs before interface assignment is critical because applying *ip vrf forwarding* removes any existing IP address from an interface.

Phase 6: PE-CE Connectivity and Routing

Customer Edge routers are connected to their respective PE interfaces. Each PE interface is assigned to the appropriate VRF and configured with an IP address. Static routes within each

VRF provide reachability to CE LANs. The BGP address family for each VRF redistributes connected and static routes, enabling VPNv4 advertisement across the backbone.

Phase 7: Requirement 2 Verification - Full Mesh

The full mesh connectivity between HR, PROD-1, and LOGISTICS is tested. VRF routing tables are examined using *show ip route vrf VRF-MESH* and *show ip bgp vpnv4 vrf VRF-MESH*. End-to-end connectivity is verified by testing all six directional paths pinging from each CE router to both remote sites.

Phase 8: Requirement 3 Verification - VPN Routing

The EXEC to R-and-D VPN is verified by examining *show ip route vrf VRF-EXEC-RD* on MAN-HO, confirming routes exist to EXEC LAN, R-and-D LAN A, and R-and-D LAN B. Ping tests confirm bidirectional connectivity before adding the IPSec overlay.

Phase 9: Requirement 1 - Secure EXEC-HR Link

The GRE over IPSec tunnel between EXEC and HR operates independently of the MPLS infrastructure, using a direct Ethernet link. IPSec ISAKMP policy 10 is configured with AES-256, SHA, and group 5 for strong security as required for sensitive HR data. The crypto map matches GRE traffic for encryption. Static routes direct LAN traffic through Tunnel0. Verification confirms IKE SA in QM_IDLE state and IPSec SA showing encrypted packet counters.

Phase 10: Requirement 3 - Selective IPSec

With VRF-EXEC-RD routing verified, IPSec is configured on EXEC and R-and-D routers for LAN A traffic only. The crypto ACL matches traffic between them, excluding LAN B. AES-128 is selected for this lower-sensitivity link, balancing security with performance. Verification confirms both LANs are reachable while only LAN A traffic triggers encryption.

Phase 11: Requirement 5 - PROD-2 Isolation

PROD-2 connects to NCL via an interface in the global routing table rather than any VRF. Since EXEC exists only within VRF-EXEC-RD, no route exchange occurs between them. Verification confirms ping failures in both directions, with PROD-2 maintaining connectivity to its local gateway.

Phase 12: Requirement 4 - SQL Server Access Control

The extended ACL SQL-SERVER-ACCESS is applied to LOGISTICS after all routing is stable. This sequencing ensures ACL hit counters reflect security enforcement rather than underlying routing problems. Testing confirms HR reaches the SQL server on port 1433, while EXEC and PROD-1 are denied.

Phase 13: Requirement 6 - DHCP Services

DHCP pools are configured last as they depend on correct interface addressing. Each CE router receives a pool matching its LAN subnet with exclusions for gateway and infrastructure addresses. The SQL server IP is explicitly excluded from the LOGISTICS pool. Verification uses *show ip dhcp pool* to confirm pool configuration.

3.2 Troubleshooting Strategy

The following outlines the troubleshooting approach used to verify correct operation and identify potential issues at each phase.

MPLS Core Diagnostics

When verifying LDP neighbours, the expected output shows three peers in Oper state on the P router. If a peer were missing, the diagnostic sequence begins with *show ip ospf neighbor* to verify OSPF adjacency, since LDP requires IGP reachability to peer loopbacks. A missing OSPF neighbour indicates layer 2 connectivity or OSPF config issues. If OSPF is operational but LDP fails, *show mpls ldp discovery* reveals whether LDP hello messages are being sent and received on the interface. Generally, the fix involves ensuring *mpls ip* is enabled on the correct interfaces.

For label distribution verification, *show mpls forwarding-table* should display labels for each PE loopback prefix. Missing labels with present routes indicate LDP session issues, while missing routes indicate OSPF advertisement problems (Cisco Systems, n.d.-a; Cisco Systems, n.d.-b).

BGP VPNv4 Diagnostics

When examining *show ip bgp summary,* peers should show established state with incrementing message counters. Peers remaining in active state indicate TCP connectivity failures to the remote loopback, resolved by verifying OSPF routes and update-source configuration. Established peers showing zero prefixes received suggest missing *send-community both* configuration, which prevents route target propagation essential for VPN route import.

The command *show ip bgp vpnv4 vrf VRF-MESH* was used to verify route exchange. Missing remote prefixes despite correct peering indicate RT mismatch. Route targets were verified using *show ip vrf detail*, ensuring import and export values align across all participating sites (Cisco Systems, n.d.-a).

IPSec Diagnostics

For Requirement 1 verification, *show crypto isakmp sa* should display QM_IDLE state indicating successful Phase 1 completion. An SA showing MM_NO_STATE suggests Phase 1 negotiation failure, commonly caused by mismatched ISAKMP policy parameters or incorrect pre-shared keys. The specific mismatch can be identified using *debug crypto isakmp* if needed (Cisco Systems, n.d.-d).

When *show crypto ipsec sa* displays zero encrypt/decrypt counters despite established ISAKMP SA, traffic is not matching the crypto ACL. This was monitored during

Requirement 3 implementation where only LAN A traffic should trigger encryption. The ACL was verified using *show access-lists* to confirm correct source and destination matching.

ACL Diagnostics

For Requirement 4, *show access-lists* displays hit counters for each line. During testing, the permit line for HR showed incrementing matches while the deny line incremented for EXEC and PROD-1 access attempts, confirming correct rule enforcement. Zero matches on expected permit lines would indicate ACL application direction issues, verified using *show ip interface* to confirm inbound vs outbound assignment (Cisco Systems, n.d.-k).

VRF Isolation Diagnostics

For Requirement 5, *show ip route* on PROD-2 confirms no route to EXEC exists. If unexpected routes appear, *show ip vrf interfaces* verifies the interface is correctly in the global table rather than a VRF. The command *show ip route vrf VRF-EXEC-RD* on NCL confirms EXEC routes remain isolated within the VRF (Cisco Systems, n.d.-a).

DHCP Diagnostics

For Requirement 6, *show ip dhcp pool* confirms pool configuration and address utilisation. If clients fail to receive addresses, *show ip dhcp binding* shows whether leases are being assigned. Empty bindings with correct pool configuration suggest layer 2 issues between client and router, while missing pools indicate the DHCP service requires verification using *show ip dhcp server statistics* (Cisco Systems, n.d.-c).

Connectivity Verification

Throughout implementation, basic connectivity was verified using ping with specific attention to response codes. Successful replies confirm reachability, while destination unreachable (U responses) suggests routing issues (if unexpected), and timeout suggests ACL filtering or interface problems. Traceroute was used to verify traffic paths, particularly confirming requirement 1 traffic traverses the GRE tunnel (192.168.100.x addresses) rather than alternative paths.

References:

Cisco Systems. (n.d.-a). How to troubleshoot the MPLS VPN. Cisco. https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13734-mpls-vpn-tsh.html

Cisco Systems. (n.d.-b). MPLS troubleshooting. Cisco. https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/12492-mpls-tsh.html

Cisco Systems. (n.d.-c). Troubleshoot DHCP in enterprise networks. Cisco. https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html

Cisco Systems. (n.d.-d). Understand and use debug commands to troubleshoot IPsec. Cisco. https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/5409-ipsec-debug-00.html

Cisco Systems. (n.d.-e). *Next generation encryption*. https://tools.cisco.com/security/center/resources/next_generation_cryptography

Cisco Systems. (n.d.-f). *Security for VPNs with IPsec configuration guide*. https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_vpnips/configuration/xe-3s/sec-sec-for-vpns-w-ipsec-xe-3s-book/sec-cfg-vpn-ipsec.html

Cisco Systems. (n.d.-g). *Dynamic multipoint IPsec VPNs (using multipoint GRE/NHRP to scale IPsec VPNs)*. https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.html

Cisco Systems. (n.d.-h). *Configure a basic MPLS VPN network*. https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html

Cisco Systems. (n.d.-i). *MPLS VPN route target rewrite*. https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/xe-16/mp-l3-vpns-xe-16-book/mpls-vpn-route-target-rewrite.html

Cisco Systems. (n.d.-j). *Introduction to Cisco IPsec technology*. https://www.cisco.com/c/en/us/td/docs/net_mgmt/vpn_solutions_center/2-0/ip_security/provisioning/guide/IPsecPG1.html

Cisco Systems. (n.d.-k). *Configure commonly used IP ACLs*. https://www.cisco.com/c/en/us/support/docs/ip/access-lists/26448-ACLsamples.html

Cisco Systems. (n.d.-l). *Configuring IP session filtering (reflexive access lists)*. https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_acl/configuration/15-mt/sec-data-acl-15-mt-book/sec-cfg-ip-filter.html

Cisco Systems. (n.d.-m). *VRF-lite software configuration guide for Cisco 1000 series connected grid routers*.

https://www.cisco.com/c/en/us/td/docs/routers/connectedgrid/cgr1000/ios/software/15_4_1_cg/vrf_cgr1000.html

Cisco Systems. (n.d.-n). *Configuring the Cisco IOS DHCP server*. https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-mt/dhcp-15-mt-book/config-dhcp-server.html